Chapter 1

Low pain vs no pain multi-core Haskells

Mustafa Aswad , Phil Trinder, Abyd Al Zain, Greg Michaelson¹, Jost Berthold²

Abstract: Multi-cores are becoming the dominant processor technology and functional languages are theoretically well suited to exploit them. In practice, however, implementing effective high-level parallel functional languages is extremely challenging.

This paper is the first programming and performance comparison of functional multi-core technologies and reports some of the first ever multi-core results for two languages. As such it reflects the growing maturity of the field by systematically evaluating four parallel Haskell implementations on a common multi-core architecture. The comparison contrasts the programming effort each language requires with the parallel performance delivered. The study uses fifteen 'typical' programs to compare a 'no pain', i.e. entirely implicit, parallel language with three 'low pain', i.e. semi-explicit languages.

The parallel Haskell implementations use different versions of GHC compiler technology, and hence the comparative performance metric is speedup which normalises against sequential performance. We ground the speedup comparisons by reporting both sequential and parallel runtimes and efficiencies for three of the languages. Our experiments focus on the number of programs improved, the absolute speedups delivered, the parallel scalability, and the program changes required to coordinate parallelism. The results are encouraging and, on occasion, surprising.

¹School of Mathematics and Computer Sciences, Heriot-Watt University,

Edinburgh, UK; {mka19, trinder, ceeatia, greg}@macs.hw.ac.uk

²Datalogisk Institut, University of Copenhagen, Denmark berthold@diku.dk

1.1 INTRODUCTION

Physical limits of semiconductor technology and improved manufacturing technologies are driving processor technology towards multi and many cores. This hardware trend has engendered much interest in functional languages, as their statelessness makes them well suited to exploit multi-core architectures, and matching interest in the functional community to exploit the new hardware, e.g. [1].

Typically a parallel functional program must not only specify the *computation* i.e. a correct and efficient algorithm, it must also specify the *coordination* e.g. how the program is partitioned, or how parts of the program are placed on processors. Most parallel functional languages incorporate high-level coordination sublanguages and a range of models have been used including data parallelism e.g. [5], *semi-explicit* models e.g. [13], coordination languages e.g. [17], and algorithmic skeletons e.g. [15]. The ultimate extreme is to make coordination entirely *implicit*, typically using either profiling as in [11] or parallel iteration as in [8]. The slogan associated with languages is 'No Pain Parallelism'.

This paper reflects the current status of multi-core functional programming by systematically comparing four parallel Haskell implementations on a common multi-core architecture (Section 1.4). We compare a 'no pain' parallel implementation Feedback Directed Implicit Parallelism (FDIP) [11], with three 'low pain', i.e. semi-explicit languages [9] (Section 1.2). The semi-explicit Haskells are Eden [13] and two implementations of Glasgow parallel Haskell (GpH) [23], namely *GpH-SMP*, an optimised shared memory implementation integrated in GHC [14], and *GpH-GUM*, a message-passing implementation designed for shared and distributed memory architectures [22] (Section 1.3).

Our performance comparisons are based on speedups, which normalise against different sequential performance. The baseline for the speedup comparisons are sequential and parallel runtimes and efficiencies for three languages (Section 1.5). Section 1.6 reports parallel performance and programming effort, focusing on the number of programs improved, speedups delivered, and program changes required to coordinate parallel evaluation. We compare the scalability, the programming effort required, and the parallel performance achieved, in each language (Section 1.7). In the concluding Section 1.8, we summarise the key results and discuss their implications. An expanded version of the paper is available as [2].

1.2 PARALLEL HASKELL LANGUAGE COMPARISON

This section outlines GpH and Eden, the two semi-explicit Haskell extensions that are compared in the remainder of the paper. The FDIP implicit approach we consider supports an unchanged Concurrent Haskell [18]. Indeed although both Eden and GpH-SMP also support concurrency i.e. multiple stateful (IO) threads, our language comparisons focuses on parallelism only. We illustrate the coordination extensions by using them to parallelise the Boyer nofib program [16]. Figure 1.1 shows the key top-level function where, in an obvious extension to the original program, the input size n is passed as a parameter.

test :: Int -> Bool
test n = all test0 (take n (repeat (Var X)))

FIGURE 1.1. Sequential Top-level Boyer function

Glasgow parallel Haskell (GpH) GpH is a modest extension of Haskell98 with parallel (par) and sequential (seq) composition as coordination primitives. *Evaluation strategies* abstract over par and seq to provide lazy, polymorphic, higher-order functions that control the evaluation degree and the parallelism of an expression [23].

```
where
                                              xs = take n (repeat (Var X))
test n m = all (&& True) res
                                              xs1 = splitAtN m xs
 where
                                              res = parallelMap (all test0) xs1
  xs = take n (repeat (Var X))
                                              parallelMap = mw np pf
  xs1 = splitAtN m xs
                                              np = noPe
  res = map (all test0) xs1
                                              pf = min 100 maxpf
         'using' parList rnf
                                              maxpf = max 2 (n 'div' (m*np*f)))
splitAtN :: Int -> [a] -> [[a]]
                                           splitAtN :: Int -> [a] -> [[a]]
splitAtN n [] = []
                                            splitAtN n [] = ... -- see on the left
splitAtN n xs = ys : splitAtN n zs
where (ys,zs) = splitAt n xs
```

FIGURE 1.2. GpH Top Boyer function

```
FIGURE 1.3. Eden Top Boyer function
```

test n m f = all (&& True) res

Figure 1.2 shows the GpH parallelization of the top-level Boyer test function, and works as follows. The input list is bound to a variable xs, and then split into chunks equivalent to physical cores and bound to xs1. Next the condition (all test0) is mapped over the chunks to give a list of intermediate results res. It is this mapping that is parallelised ('using' parList rnf). The final stage is to combine the intermediate results all (&& True) res.

The parallelization illustrates some interesting points. In this program, just 1 of the 52 functions in the 300 line program changes. This is the case for many, but not all, programs (exceptions are Sphere and Hidden where parallelism is introduced in more than one function). The parallel paradigm is chunked data parallelism. That is, the parallelism is determined by the underlying data structure, and to obtain suitable thread granularity, the program has been changed to aggregate the input. In other programs it is possible to introduce parallelism without changing the algorithmic or computational part of the program, e.g. [12].

Eden Eden [13] extends Haskell with syntactic constructs to define and instantiate processes explicitly. In contrast with the other languages, such direct Eden programming exposes parallel tasks at the language level, and requires the programmer to manage them using the control mechanisms provided in the language. In practice however, Eden provides libraries of skeletons [13, 3] and many programs, including all of the nofib suite here, can be parallelised using them.

Eden supports a distributed memory parallel paradigm. That is, processes share no values, and communicate only by messages. It might be thought that such a paradigm would not be suitable for parallelism on shared-memory multi-core architectures; however recent results have shown good performance [1], as indeed do the results in Sections 1.5, 1.6, and 1.7. We return to this issue in Section 1.8.

Description	FDIP	GpH	Eden
Classification	Implicit	Semi-explicit	Semi-explicit
Evaluation Order	NormalOrder	Normal/ Mixed	Mixed
Methodology	FDIP Tools	Evaluation Strategies	Direct or Skeletons
Process Model	Speculative Threads	Optional	Explicit Processes
& Creation			Mandatory Creation
Thread Placement	Implicit	Implicit & Dynamic	Implicit & Static
Communication	Implicit	Implicit	Implicit & Explicit
Channels			

 TABLE 1.1.
 Language-level Comparison of Parallel Haskells

Figure 1.2 shows the Eden skeleton-based parallelization of the top-level Boyer test function, and works as follows. As before, the input list is chunked into xs1 and the intermediate results combined by all (&& True) res in the final stage. For Eden, the mapping of (all test0) over the chunks is parallelised using a master worker skeleton parameterised by the number of cores and the number of tasks to prefetch: mw np pf. The *f* parameter specifies the number of tasks. The average tasks per worker is list length (*n*) divided by chunk size (*m*) and no. of workers (*np*), that is $\lfloor \frac{n}{m \cdot np \cdot f} \rfloor$. As in GpH the paradigm is chunked data parallelism, and just one out of 52 functions has been parallelised, although this time an algorithmic skeleton is used.

Language Comparison Table 1.1 summarises the language level differences in coordination specification in the three parallel Haskells. Much of the table summarises aspects outlined above. However a key distinction between the languages is that while FDIP preserves normal order evaluation of pure expressions, GpH may not, and Eden does not. GpH preserves normal order evaluation if every evaluation strategy added is no more strict than the embedding function. However it is often useful to be more strict, e.g. speculatively evaluating expressions in the anticipation that they will be used. While Eden processes preserve some normal order evaluation internally, e.g. of expressions within the body of a process, processes will always strictly evaluate their arguments and results.

As an entirely implicit language, FDIP has the highest level of coordination abstraction, GpH an intermediate level and Eden the lowest. That is, Eden is most explicit about coordination behaviour, but as we shall see in Section 1.7, the use of appropriate skeletons can raise the level of abstraction.

1.3 PARALLEL HASKELL IMPLEMENTATION COMPARISON

All of the parallel Haskells perform parallel graph reduction [19] and support high-level coordination, relying on sophisticated implementations to effectively manage a vast array of low-level coordination issues, typically including task placement, communication, synchronisation, and storage management.

Feedback Directed Implicit Parallelism (FDIP) Parallelism is introduced and controlled in FDIP in a four stage process [11] as follows. Firstly, an example execution of the program is profiled. Secondly, the profile trace is analysed as a dependency graph of computations to identify useful sources of parallelism. Given the large number of potential computations *thunks* in almost any Haskell

program, the challenge is to identify thunks that are simultaneously independent of other thunks, demanded by the program, and with large thread granularity. The third stage is to recompile the program to introduce parallelism automatically at the identified program sites. Finally, sophisticated mechanisms are introduced into the runtime system to manage the threads introduced at these sites. These include treating the parallel threads as speculative, and managing load with work stealing.

GpH-SMP Since 2004 the Glasgow Haskell Compiler (GHC) has supported a shared-memory implementation of GpH. The shared memory implementation is evolving rapidly [14], and the precise version we describe here and measure in later Sections is the head version GHC 6.11 from June 2009, denoted *GpH-SMP* in the remainder of the paper. The GHC runtime system implements Concurrent Haskell threads using *capabilities* and a system of lightweight threads multiplexed onto a small number of heavyweight OS threads to achieve real parallelism on a multiprocessor, while still keeping overheads of concurrency low [10]. GHC 6.11 supports both parallel and sequential garbage collection, and the measurements in the following sections use the former. In this scheme, when memory is exhausted all cores cease reduction and perform garbage collection in parallel.

GUM Implementation of GpH Graph-reduction on a Unified Machine-model (GUM) is a portable, parallel runtime environment for GpH [22]. As the name suggests, GUM is designed for both shared and distributed memory architectures. It implements a Distributed Shared Memory (DSM) model of parallel graph reduction on a distributed, but virtually shared, graph. Graph segments are communicated in a message passing architecture, using standard communication libraries like PVM [20] to provide an architecture neutral and portable runtime environment.

Eden Implementation The Eden implementation extends GHC making a few changes to the front-end, and major modifications to the runtime environment [4]. When run in parallel, each PE runs a sequential copy of the GHC runtime system. Multiple PEs communicate by message-passing, and the communication layer has been designed to allow plug-in replacement of different message-passing libraries. Typically, it uses either PVM or MPI libraries.

Implementation Comparison Table 1.2 summarises the implementation level differences between the four parallel Haskells. While an arbitrary number of Eden processes can be dynamically created, each process is mandatory. In contrast, the other implementations support dynamic techniques including thread subsumption, sparking, and the creation of optional or speculative threads. Eden uses eager work distribution: newly created processes are pushed out to available PEs, while the other implementations are lazy with idle PEs stealing work (thunks). FDIP and GpH-GUM are careful not to duplicate work by evaluating the same thing more than once, but work may be duplicated in Eden and, sometimes, in GpH-SMP.

A key distinction between the implementations is the heap model: while FDIP and GpH-SMP have shared heaps, GUM maintains a virtual shared heap, and Eden uses distributed independent heaps, both interacting via message passing. Mes-

Description	FDIP	GpH-SMP	GpH-GUM	Eden
GHC Version	GHC 6.6	GHC 6.11	GHC 4.06	GHC 6.8
Evaluation	Par. Graph	Par. Graph	Par. Graph	Par. Graph
Model	Reduction	Reduction	Reduction	Reduction
Granularity Ctrl	Dynamic	Dynamic	Dynamic	Static
Synchronisation	Thunk	Thunk	Thunk	Channel
Unit	Locking	Locking	Locking	Locking
Work	Work	Work	Work	Dynamic
Distribution	Stealing	Pushing	Stealing	Process
		& Stealing		Placement
Work Duplication	Not Possible	Possible	Not Possible	Possible
Неар	Shared	Shared	Virtual	Distr.
	Неар	Неар	Shared Heap	Неар
Garbage	Dependent,	Dependent,	Independent,	Independent,
Collection	Sequential	Parallel	Parallel	Parallel

TABLE 1.2. Implementation-level Comparison of Parallel Haskells

sage passing is essential for distributed systems, but initially seems enormously expensive compared with shared memory access. That is, computation graphs must be serialised into, and deserialised from, messages, and potentially expensive message-passing functions invoked. However the independent heaps maintained by GUM and Eden convey four significant advantages for shared-memory systems like multi-cores. Firstly, while the cores in shared heap implementations like FDIP and GpH-SMP must synchronise to garbage collect, GUM and Eden cores can collect independently and hence in parallel. Secondly, synchronisation buffers. Thirdly, synchronisation granularity is often large, i.e. on large messages, rather than on individual thunks or memory locations. Finally cache coherency issues are reduced as tasks do not share caches [1]. We discuss the performance implications of the heap designs further in Section 1.8.

Although both FDIP and GpH-SMP use dependent stop-the-world garbage collection, such a design is not inherent. An implementation that maintains some form of thread-private heap, e.g. [6], would enable independent garbage collection and offer many of the advantages outlined above, without incurring the high communication costs of message passing.

1.4 EXPERIMENT DESIGN

We compare the performance of the four parallel Haskells using fifteen programs from the 'real' and 'spectral' sections of the nofib benchmark suite [16]. The 'real' and 'spectral' sections of the nofib suite are carefully designed to be representative of small Haskell programs, i.e. around 300 source lines of code. The programs are a substantial subset of the twenty multi-core benchmarks used in [11] that are in turn carefully selected to be representative. Of the five programs *not* measured, two are not nofib benchmarks, and three (cacheprof, calendar and fibheaps) are too small to benefit from parallel execution, i.e. where the input cannot be sized to give a runtime of 3s on current hardware. Crucially, other than

to exclude short programs, the programs are not selected *a priori* for having obvious parallel structure. Hence our results reflect the multi-core performance that might be expected for a set of 'typical' small Haskell programs. To parallelise the programs in Eden and GpH the programs are first time and space profiled to identify computationally expensive functions, and these are parallelised. A variety of parallelizations are investigated for each program and the best selected. The same GpH program is evaluated under GpH-SMP and GpH-GUM, and the Eden program introduces an appropriate skeleton. Example GpH and Eden parallelizations of the Boyer benchmark are discussed in 1.2.

All programs are measured on the same input, and with the same heap size. We follow the common practice of increasing input size in many cases to match improvements in processor technology since the benchmarks were established in 1992. The best parallel performance is reported for each system. For Eden, GpH-SMP and GpH-GUM, the best performance is obtained on eight cores, but for FDIP it is obtained on four cores as discussed in Section 1.7. Parallel runtimes are variable and to ameliorate these effects the measurements are based on the median runtimes from three executions.

The implementations are all based on the GHC compiler, but use different versions of it. The FDIP approach uses GHC 6.6, GpH-SMP uses GHC 6.11, GpH-GUM uses GHC 4.06, and Eden uses GHC 6.8. GHC has been continuously developed to maturity over the years, and typically the sequential execution time of programs is improved by later versions of the compiler. To address the issue of varying sequential performance, the primary comparative measure is *absolute speedup*, i.e. relative to the corresponding optimised sequential GHC compiler, e.g. Eden speedups are relative to GHC 6.8. The absolute speedups are grounded by comparative runtime measurements in Section 1.5.

The programs are all measured on common multi-core architectures, namely eight core machines comprising two quad-cores. The GpH-SMP, GpH and Eden measurements are for Intel Xeon 5410 cores running at 2.33GHz, with a 1998 MHz front-side bus 6144 KB and 8GB RAM running under Linux Fedora 7. The FDIP measurements are for Intel Xeon X5350 running at 2.66GHz with 4GB RAM running under Windows Server 2003 R2 x64 service pack 2.

1.5 RUNTIME COMPARISON

As the parallel implementations use different versions of the GHC compiler (Section 1.4), this section provides a baseline for the speedup measurements in the following sections by comparing the runtimes and efficiencies of the GpH-SMP, GpH-GUM and Eden parallel implementations on 1, 2, 3, 4, 6, and 8 cores. FDIP is excluded as an implementation is not available.

Single Core Results Table 1.3 summarises the single core runtimes of four nofib programs that deliver good speedup. To facilitate comparison, the inputs for the programs are sized to give sequential GHC 6.11 runtimes of approximately 35s. Columns 2–4 of the table report optimised sequential runtimes for the compiler instances extended by the parallel Haskell implementations, and these form the basis for the absolute speedup calculations in the remainder of the paper.

	Sequential			1 Core			8 Core		
Program	GHC	GHC	GHC	GpH	GpH	Eden	GpH	GpH	Eden
	6.11	4.06	6.8	SMP	GUM		SMP	GUM	
Boyer	34.1	49.3	36.7	35.0	77.52	37.1	10.0	14.1	10.1
Clausify	26.9	51.2	29.1	33.1	78.7	29.3	3.9	11.5	3.9
Fft2	38.1	75.7	48.6	38.2	80.9	49.2	13.3	45.8	17.7
Rewrite	37.1	68.1	46.8	39.0	94.9	52.05	6.8	26.9	9.9
Mean	34.1	61.1	40.3	36.3	83.0	41.9	8.5	24.6	10.7

TABLE 1.3. Sequential and Parallel Runtime Comparison (seconds).



FIGURE 1.4. Runtime Comparison of Parallel Haskells (Boyer)

Columns 5–7 of the table report the 1 core parallel runtimes for each implementation. Columns 8–10 of the table report the 8 cores parallel runtimes for each implementation. We make the following observations.

1: The sequential runtimes vary by as much as a factor of 2.0: Fft2 under GHC 6.11 takes 38.1s, and under GHC 4.06 takes 75.7s, but typical variation is less.

2: The mean sequential runtimes show that GHC 4.06 is the slowest on a single core, and 1.8 (61.1/34.1) times slower than GHC 6.11. This reflects recent GHC performance improvements. GHC 6.8 is 18% (40.3/34.1) slower than GHC 6.11. Longer runtimes for GHC 4.06 and GHC 6.8 give GpH-GUM, and to a lesser extent Eden, an advantage in the following speedup measurements as the compute time is relatively large compared with communication time.

3: When executed on a single processor parallel language implementations introduce overheads compared with sequential execution, e.g. launching a single virtual PE [22]. This *sequential efficiency* is a function of both the parallel program and the architecture, and represents this parallel overhead. Comparing columns 2 and 5, 3 and 6, and 4 and 7 of Table 1.3 shows that this expectation is met for all languages with mean sequential efficiencies of 74% (61.1/83.0) for GpH-GUM, 96% (40.3/41.9) for Eden, and 94%(34.1/36.3) for GpH-SMP.

Eight Core Results Table 1.3 (columns 8–10) summarises the runtimes of the same four programs on eight cores. We make the following observations.

1: On eight cores the variation in runtimes is at most a factor of 4.1 (26.9/6.5), between GpH-SMP and GpH-GUM Rewrite, but is typically rather less.

2: The mean eight core runtimes show that, for this collection of programs, GpH-SMP remains fastest, Eden is just 26% (10.7/8.5) slower, and GpH-GUM slowest

by a factor of 2.9 (24.6/8.5).



FIGURE 1.5. Absolute Speedup Comparison of Parallel Haskells (Boyer)

Runtime and Speedup Graphs Figures 1.4 and 1.5 compare the runtimes and absolute speedups of Boyer program from Table 1.3. The program is measured on 1, 2, 3, 4, 6, and 8 cores. We make the following observations.

1: The runtime curves are broadly similar for all implementations. For GpH-SMP and Eden the curves are very similar, and while the Eden is a little slower on one core, the eight core results are very similar.

2: Three implementations *scale*, i.e. the runtimes fall as cores are added. The only exceptions are between two and four cores under GpH-GUM and between four and eight cores under GpH-SMP and Eden. This is in marked contrast to FDIP, where the best performance may be achieved under two, three, or four cores [11], and we shall return to this point in Section 1.7.

3: Reflecting the runtime curves, the speedup curves for the program are broadly similar, and for GpH-SMP and Eden they are very similar.

4: The speedup on a single core reflects the sequential efficiencies of the implementations.

1.6 PROGRAMMING EFFORT AND PERFORMANCE RESULTS

This section investigates the parallel performance of the four parallel Haskells in conjunction with the programming effort required to achieve that performance. Parallel performance is measured as *absolute* speedup and programming effort as logical source lines of code (SLOC) [7], both as an absolute number and as a percentage of program length. For our purposes SLOC has the advantages of simplicity and relatively wide use. We also record the parallel paradigm applied in GpH and Eden. The absolute speedups achieved for all fifteen programs in the four languages are depicted in Figure 1.7.

FDIP Multi-Core Performance FDIP is entirely implicit, and so no programmer effort is expended other than in profiling and using a special compiler. Simi-

Program Name	1 8	Speedup	Lines Of Code
Hidden		1.82	316
sort1		1.3	113
sort2		1.3	124
Atom		1.27	57
Simple		1.27	1053
Mean		1.5	
TABLE	1.4.	FDIP	Programs
Improved			

Program	Speedup	Lines	Lines	%	Paradigm	
Name		Code	Chgd	Chgd	_	
Clausify	6.8	101	6	6	Chunked Data	
					Parallelism	
Primetest	5.9	112	15	13	Chunked Data	
					Parallelism	
Rewrite	5.5	408	14	3	Chunked Data	
					Parallelism	
Sphere	3.8	332	12	4	Nested Data	
					Parallelism	
Hidden	3.5	316	6	2	Nested Data	
					Parallelism	
Boyer	3.4	295	9	3	Chunked Data	
					Parallelism	
Fft2	2.9	705	13	2	Data Parallelism	
Mean	4.5		10.7	4.7		
TABLE 1.5. GpH-SMP Progs Improved						

larly the programmer does not need to identify and apply some parallel paradigm. The FDIP performance results reported in this paper are based on the ICFP'07 paper [11], augmented with some additional results from the authors. Where the other parallel Haskells are measured on eight cores, FDIP performs better on four cores than on eight and hence Table 1.4 follows [11] in reporting the programs

the programs, with a mean speedup of 1.5, and maximum speedup of 1.8 Automatically extracting good parallel performance is acknowledged to be a challenging problem. However, some of the reasons for the relatively poor performance of FDIP are due to the implementation being immature compared with the other systems and having some known technical problems [11]. Specifically, the simulation ignores several crucial aspects of parallel coordination, namely contention within the GHC runtime system; the overheads of the shim lock implementation; and finally the overheads of sparking work and the cache effects of moving data from a sparking core to one running work speculatively.

that are improved on four cores. It shows that FDIP speeds up only 20% (3/15) of

GpH-SMP Multi-Core Performance Table 1.5 reports the programming effort and parallel performance of programs improved by GpH-SMP on eight cores. As a semi-explicit parallel language, GpH requires the programmer to identify a suitable parallel paradigm and introduce evaluation strategies to apply it. Introducing the parallelism requires changing an average of just 10.7 lines in each program, i.e. 4.7% of the code.

The table shows that GpH-SMP improves nearly half of the programs, i.e. 46% (7/15). The mean speedup is 4.5, with a best speedup of 6.8 for Clausify. It is impressive that three of the programs achieve speedups of four or more on eight cores, i.e. a parallel efficiency of 50% or more.

GpH-GUM Multi-Core Performance Table 1.6 reports the programming effort and parallel performance of programs improved by GpH-GUM on eight cores. Only twelve of the fifteen programs are attempted for GpH-GUM as Compress, Hidden and Primetest import modules not available in GHC 4.06.

As before, GpH requires the programmer to identify a suitable parallel paradigm and apply it. Introducing the parallelism requires changing an average of just eleven lines of each of these programs, i.e. 3.6% of the code. The table shows

Program	Speed-	Lines	Lines	%	Paradigm
Name	up	Code	Chgd	Chgd	_
Clausify	4.5	101	6	6	Chunked Data
					Parallelism
Boyer	3.5	295	9	3	Chunked Data
					Parallelism
Rewrite	2.5	408	14	3	Chunked Data
Sphere	1.8	332	12	4	Nested Data
					Parallelism
Fft2	1.7	705	13	2	Data Parallelism
Mean	3.1		11	3.6	

Program	Speed-	Lines	Lines	%	Paradigm
Name	up	Code	Chgd	Chgd	_
Clausify	6.2	101	7	7	Data Parallelism
Rewrite	4.7	408	15	4	Chunked Data
					Parallelism
Fft2	3.7	705	11	2	Data Parallelism
Boyer	3.6	295	14	5	Chunked Data
					Parallelism
Compress	1.6	109	3	2	Data Parallelism
Sphere	1.5	332	7	2	Data Parallelism
Mean	3.6		10	3.6	

TABLE 1.6. GpH-GUM Progs Improved TABLE 1.7. Eden Progs Improved

Description	FDIP*	GpH-	GpH-	Eden			
		SMP	GUM				
No. Programs Measured	15	15	12	15			
No. Programs Improved	3	7	5	6			
% Programs Improved	20%	46%	42%	40%			
No. Lines Changed	0	11	11	10			
% Code Changed	0	4.7	3.6	3.6			
Mean Speedup	1.5*	4.5	3.1	3.6			
* Performance on four Cores							

TABLE 1.8. Comparative Multi-Core Performance Summary

that GpH-GUM improves 42% (5/12) of the programs. The mean speedup is 3.1, with a best speedup of 4.5 for Clausify.

Eden Multi-Core Performance Table 1.7 reports the programming effort and parallel performance of programs improved by Eden on eight cores. Eden requires the programmer to identify a suitable parallel paradigm and introduce an appropriately parameterised algorithmic skeleton to exploit it. This set of programs all use the master-worker skeleton discussed in Section 1.2, but some do so directly, while others like Boyer and Rewrite chunk the input to improve thread granularity. Introducing the parallel coordination requires changing an average of just ten lines in each program, again just 3.6% of the program text.

The table shows that Eden improves a slightly smaller fraction, i.e. 40% (6/15), of the programs than GpH-SMP and GpH-GUM. The maximum speedup is similar to GpH-SMP (6.2 vs 6.8), the mean speedup is slightly less (3.6 vs 4.7). Four of the programs achieve speedups of 3.6 or more on eight cores, i.e. a parallel efficiency of 45% or more.

1.7 COMPARATIVE STUDY

This section compares the best parallel performance of the four Haskell languages and the programming effort required to achieve that performance. Table 1.8 summarises the key metrics from section 1.6.

Programming Effort Comparison As a purely implicit approach, FDIP requires minimal programmer effort, simply the execution of a profiling run. In contrast GpH and Eden both require programmer effort to time-profile the program, to



FIGURE 1.6. Performance Scalability of Parallel Haskells on Four Cores.

insert evaluation strategies or skeletons, and to tune the parallel performance. Tables 1.5, 1.6, and 1.7 show that the scale of the program changes is on average small in both absolute and relative terms, e.g. representing just eleven lines or 4.3% of the program text in both languages. We conclude that, for these relatively simple programs, using existing Eden skeletons represents a similar level of coordination abstraction to evaluation strategies in GpH.

The results also illustrate that in both GpH and Eden some programs are easier to parallelise than others. That is, the scale of program changes induced by parallelization may vary significantly in both absolute and relative terms. For example, Table 1.5 shows that in GpH the number of lines changed may vary from three to fifteen, and the percentage of program text may vary from 1% to 13%. Similarly, Table 1.7 shows that in Eden the number of lines changed may vary from three to fifteen, and the percentage of program text may vary from 2% to 7%.

The parallel paradigms used in the improved programs are all forms of data parallelism, sometimes combined with *chunking* to increase thread granularity, or *nesting* to introduce additional parallelism. Although the parallelization changes are small, the source lines of code metric does not reflect the programmer effort expended on the key intellectual challenge, namely understanding the computational structure of a program written by another programmer.

Scalability A key property of a parallel implementation is *scalability*, i.e. whether performance increases as processing elements are added. We have already seen the scalability of the GpH-SMP, GpH-GUM and Eden implementations up to eight cores in the discussion of Figure 1.4.

Figure 1.6 provides a more detailed analysis for three programs (Boyer, FFT2 and Simple) in each language on 1, 2, 3 or 4 cores. Each program gives good performance on at least one implementation. The figure shows that in GpH-SMP, GpH-GUM and Eden the performance of programs, i.e. speedup, for Boyer and FFT2 improves steadily as cores are added. In contrast FDIP delivers the best speedup for Simple on three cores. This is not an isolated result: the 5 programs delivering speedups under FDIP reported in [11] deliver maximum speedup twice on three cores, and three times on four cores.



FIGURE 1.7. Performance Comparison of Parallel Haskells (8 cores)

Furthermore, FDIP ceases to scale beyond four cores [21], and this is illustrated by the four core performances of Boyer, Simple and FFT2 in Figure 1.6, which are uniformly better than the 8 core performances reported in Figure 1.7. The reasons for this have not been established, but are likely to be either lock contention or low-level memory effects, e.g. disrupting caches when transferring threads between cores.

Performance Comparison A complete comparison of the speedups achieved on eight cores for all fifteen programs in the four languages is depicted in Figure 1.7. The performance price of FDIP's purely implicit approach is high: it improves the fewest number of programs: three out of fifteen on four cores (Table 1.4), and two out of fifteen on eight cores (Figure 1.7). Moreover the mean and maximum speedup are both relatively small at 1.5 and 1.8 respectively on four cores. However, a mean speedup of 1.5 on four cores shows parallel efficiency approaching that of the semi-explicit implementations, i.e. speedups of approximately 3.5 on eight cores. FDIP parallelism scales both irregularly, and only to four cores (Section 1.7).

GpH-SMP delivers the best performance, improving 7/15 programs with mean speedup of 4.5 and maximum speedup of 6.8. Eden delivers the next best performance, improving 6/15 programs with mean speedup of 3.6 and maximum speedup of 6.2. GpH-GUM is the worst of the semi-explicit languages, improving 5/12 programs with mean speedup of 3.1 and maximum speedup of 4.5. We analyse the implications of these relative performances in section 1.8.

1.8 CONCLUSION

Summary The preceding sections report the first comparison of functional multicore technologies and are some of the first ever GpH-GUM and GpH-SMP multicore results. We contrast a 'no pain' approach with three 'low pain' approaches, and start by outlining and comparing the approaches at both language and implementation levels. We present the design of an experiment using fifteen programs carefully selected to reflect the multi-core performance that might be expected for a typical set of Haskell programs.

To ground the speedup comparisons we report sequential and parallel runtimes and efficiencies for three of the languages. Although the parallel Haskell implementations all use GHC, they each use a different version, and hence expose different relative speedups. This should be taken into account when looking at our primary performance metric (absolute speedups). We find that sequential runtimes vary by as much as a factor of 2.9, and eight core runtimes by as much as a factor of 4.0 (26.9/6.8). On a single core GpH-SMP is fastest and GpH-GUM slowest, and sequential efficiencies vary between 74% and 95%. Finally runtime and speedup graphs show that GpH-SMP, GpH-GUM and Eden parallel performance scales.

We report detailed parallel performance and programming effort studies, and make a comparative study with the following key results.

1: FDIP's purely implicit approach requires minimal programmer effort. In contrast GpH and Eden both require programmer effort to understand the program's computational structure, to profile it, to insert parallel coordination, and to tune the parallel performance. As the languages provide high levels of coordination abstraction the program changes are small, on average no more than 4.3% of the program text in both languages. We conclude that Eden skeletons represent a similar high level of coordination abstraction to evaluation strategies in GpH.

2: While GpH-SMP, GpH-GUM and Eden all scale consistently up to eight cores, FDIP does not scale beyond four cores and may deliver best performance on three or four cores.

3: The performance price of FDIP's purely implicit approach is high: it improves the fewest number of programs (just three out of fifteen) and the mean and maximum speedup are both relatively small at 1.5 and 1.8 respectively on four cores.

4: All three semi-explicit approaches improve approximately half of the programs, and the performance of GpH-SMP exceeds Eden which in turn exceeds GpH-GUM.

Discussion As multi-cores become the dominant processor technology it is crucial that functional languages realise their theoretical potential to exploit them effectively. Our study reflects some of the technologies emerging to do so, namely four multi-core Haskell implementations, and the results have a number of implications for the field.

It is clear that purely implicit parallelism remains an elusive goal. The FDIP approach speeds up fewer programs, with smaller speedups, and does not scale well. While it is not clear that the scaling issues with FDIP are fundamental, the move towards many cores will make scalability a crucial property for languages and implementations.

It might be seen as discouraging that, even in the low pain languages, only half of the programs deliver absolute speedups, and that the mean parallel efficiencies are only around 45% (Tables 1.5, 1.6, and 1.7). However, recall that these programs were neither designed to be parallel, nor selected for their inherent parallelism. While some algorithms will remain inherently sequential, it is likely that with thoughtful design a far higher percentage of programs can be ef-

fectively parallelised. Moreover, the implementations are evolving fast and we can expect greater parallel efficiencies in the near future.

Interestingly the multi-core performance delivered by Eden remains comparable with that of GpH-SMP, e.g. a mean speedup within 30%, and a maximum speedup within 10%. This is surprising as while the GpH-SMP implementation is designed for multi-cores and communicates/synchronises via shared memory access, the Eden implementation is designed for distributed memory architectures, and pays a massive communication/synchronisation overhead for message passing. Passing a message entails serialising heap, calling expensive communication libraries, and deserialising heap. We argue that the key reason for the relatively good multi-core performance of Eden is the maintenance of *independent heaps*, and Section 1.3 discusses the four significant advantages of independent heaps for multi-cores. Furthermore, we speculate that as multi-core scale to many cores the advantages of independent heaps will be greatly magnified, and that some form of thread-private heap, e.g. [6], will be essential.

There are many encouraging signs for multi-core functional languages. The GpH and Eden semi-explicit approaches deliver effective high-level coordination, and hence require very small program changes, and perhaps only half a working day to introduce and tune the parallelism for a known program. The fact that there are four multi-core Haskells to compare reflects the level of interest in addressing the challenges. The implementations all have considerable room for improvement, e.g. the GpH-GUM and Eden naive distributed memory implementations could be dramatically improved by shared-memory communication libraries. Another promising line of future work is to integrate distributed and shared-memory implementations to better exploit the increasingly ubiquitous clusters of multi-core architectures.

Acknowledgements Thanks to Rita Loogen, Tim Harris, and Simon Marlow for constructive feedback. This research is supported by European Union Framework 6 grant RII3-CT-2005-026133 SCIEnce.

REFERENCES

- [1] A. Al Zain, J. Berthold, K. Hammond, P. Trinder, G. Michaelson, and M. Aswad. Low-Pain, High-Gain Multicore Programming in Haskell: Coordinating Irregular Symbolic Computations on MultiCore Architectures. In Workshop on Declarative Aspects of Multicore Programming (DAMP'09). ACM Press, 2009.
- [2] M. Aswad, P. W. Trinder, A. Al Azain, G. Michaelson, and J. Berthold. Low Pain vs No Pain Multi-core Haskells (Full Report). Technical Report 0068, http://www.macs.hw.ac.uk/~mka19/papertec.pdf, School of Math. Comp. Sciences Heriot-Watt Univ., Edinburgh, UK, Jun 2009.
- [3] J. Berthold, M. Dieterle, R. Loogen, and S. Priebe. Hierarchical master-worker skeletons. In PADL'08 — Practical Aspects of Declarative Languages, pages 248–264, 2008.
- [4] J. Berthold and R. Loogen. Parallel coordination made explicit in a functional setting. In Z. Horváth and V. Zsók, editors, 18th Intl. Symposium on the Implementation of Functional Languages (IFL 2006), Springer LNCS 4449, Budapest, Hungary, 2007.

- [5] M. M. T. Chakravarty, R. Leshchinskiy, S. Jones, G. Keller, and S. Marlow. Data Parallel Haskell: a status report. In *Workshop on Declarative Aspects of Multicore Programming (DAMP'07)*. ACM Press, 2007.
- [6] T. Domani, G. Goldshtein, E. K. Kolodner, E. Lewis, E. Petrank, and D. Sheinwald. Thread-local heaps for java. In SIGPLAN Not, pages 76–87. ACM Press, 2002.
- [7] N. E. Fenton and S. L. Pfleeger. Software Metrics: A Rigorous and Practical Approach. PWS, 1998.
- [8] C. Grelck and S. B. Scholz. SaC from High-Level Programming with Arrays to Efficient Parallel Execution. *Parallel Processing Letters*, 13(3):401–412, 2003.
- [9] K. Hammond and G. Michaelson, editors. *Research Directions in Parallel Functional Programming*. Springer, London, 1999.
- [10] T. Harris, S. Marlow, and S. Jones. Haskell on a Shared-Memory Multiprocessor. In Proc. ACM Haskell'05, pages 49–61, New York, NY, USA, 2005. ACM.
- [11] T. Harris and S. Singh. Feedback directed implicit parallelism. SIGPLAN Not., 42(9):251–264, 2007.
- [12] H.-W. Loidl, P. Trinder, K. Hammond, S. Junaidu, R. Morgan, and S. Peyton Jones. Engineering Parallel Symbolic Programs in GPH. *Concurrency — Practice and Experience*, 11(12):701–752, October 1999.
- [13] R. Loogen, Y. Ortega-Mallén, and R. Pena-Marí. Parallel Functional Programming in Eden. *Functional Programming*, 2005.
- [14] S. Marlow, S. Peyton Jones, and S. Singh. Runtime Support for Multicore Haskell. In *ICFP'09*, 2009. (to appear).
- [15] G. Michaelson, H. S., N. Scaife, and P. Bristow. A Parallel SML compiler based on algorithmic skeletons. *Journal of Functional Programming*, 15(4):615–650, 2005.
- [16] W. Partain. The nofib Benchmark Suite of Haskell Programs. In *Glasgow Work-shop on Functional Programming*, Workshops in Computing, pages 195–202, Ayr, Scotland, 1992. Springer-Verlag.
- [17] J. Peterson, V. Trifonov, and A. Serjantov. Parallel Functional Reactive Programming. In PADL'00 — Practical Aspects of Declarative Languages, LNCS 1753, pages 16– 31. Springer-Verlag, 2000.
- [18] S. Peyton Jones, A. Gordon, and S. Finne. Concurrent Haskell. In POPL'96 Symposium on Principles of Programming Languages, pages 295–308, St Petersburg, Florida, Jan. 1996. ACM.
- [19] S. Peyton Jones and D. Lester. *Implementation of Functional Programming Languages*. Prentice Hall, 1987.
- [20] Parallel Virtual Machine Reference Manual. University of Tennessee, Aug 1993.
- [21] S. Singh. Private Communication about FDIP Performance, 11 2008.
- [22] P. W. Trinder, J. Hammond, J. Mattson, A. Partridge, and S. L. Peyton Jones. GUM: a portable parallel implementation of Haskell. In *Proc. ACM PLDI'96*, pages 79–88, New York, NY, USA, 1996. ACM Press.
- [23] P. W. Trinder, K. Hammond, H. W. Loidl, and P. Jones. Algorithm + Strategy = Parallelism. *Journal of Functional Programming*, 8:23–60, 1998.