

Chapter 1

Designing a Generic Graph Library using ML Functors

Sylvain Conchon¹, Jean-Christophe Filliâtre¹, Julien Signoles²

Abstract: This paper details the design and implementation of OCAMLGRAPH, a highly generic graph library for the programming language OCAML. This library features a large set of graph data structures—directed or undirected, with or without labels on vertices and edges, as persistent or mutable data structures, etc.—and a large set of graph algorithms. Algorithms are written independently from graph data structures, which allows combining user data structure (resp. algorithm) with OCAMLGRAPH algorithm (resp. data structure). Genericity is obtained through massive use of the OCAML module system and its functions, the so-called *functors*.

1.1 INTRODUCTION

Finding a graph library for one’s favorite programming language is usually easy. But applying the provided algorithms to one’s own graph data structure or building undirected persistent graphs with vertices and edges labeled with data other than integers is likely to be more difficult. Figure 1.1 quickly compares several graph libraries according to the following criteria: number of graph data structures; purely applicative or imperative nature of the structures; and ability to apply the provided algorithms to a user-defined graph data structure. As one can notice, none of these libraries gives full satisfaction. This paper introduces OCAMLGRAPH³, a highly generic graph library for the programming language OCAML [7], which intends to fulfill all these criteria.

¹LRI, Univ Paris-Sud, CNRS, INRIA Futurs (ProVal), 91893 Orsay, France;
Email: {Sylvain.Conchon, Jean-Christophe.Filliatre}@lri.fr

²CEA-LIST, Laboratoire Sûreté des Logiciels, 91191 Gif-sur-Yvette cedex, France;
Email: Julien.Signoles@cea.fr

³OCAMLGRAPH is an open source library available at <http://ocamlgraph.lri.fr/>.

	language	graph data structures	imperative / persistent ⁴	generic algorithms
GTL [5]	C++	1	I	no
LEDA [13]	C++	2	I	no
BGL [2]	C++	2	I	yes
JDSL [4]	Java	1	I	yes
FGL [3, 9]	Haskell	1	P	yes
MLRisc [6]	SML	1	I	no
Baire [1] ⁵	OCAML	8	P/I	meaningless

FIGURE 1.1. Comparison with other graph libraries

OCAMLGRAPH introduces genericity at two levels. First, OCAMLGRAPH does not provide a single data structure for graphs but many of them, enumerating all possible variations (19 altogether)—directed or undirected graphs, persistent or mutable data structures, user-defined labels on vertices or edges, etc.—under a common interface. Secondly, OCAMLGRAPH provides a large set of graph algorithms that are written independently from the underlying graph data structure. These can then be applied naturally to the data structures provided by OCAMLGRAPH itself and also on user-defined data structures as soon as these implement a minimal set of functionalities.

Without proper parameterization, such a large set of variants may easily result in unmanageable code. We avoid this pitfall using the OCAML module system [12], which appears to be the tool of choice for this kind of meta-programming. The genericity of OCAMLGRAPH is indeed achieved through a massive use of OCAML functors. On one hand, they are used to avoid code duplication between the many variations of graph data structures, which is mandatory here due to the high number of similar but different implementations. On the other hand, they are used to write graph algorithms independently from the underlying graph data structure, with as much genericity as possible while keeping efficiency in mind.

This paper is organized as follows. Section 1.2 gives an overview of OCAML module system. Section 1.3 demonstrates the use of OCAMLGRAPH through an example. Section 1.4 exposes the design of the common interface for all graph data structures and explains how the code is shared among various implementations. Section 1.5 describes the algorithms provided in OCAMLGRAPH and how genericity is obtained with respect to the graph data structure. Finally Section 1.6 presents some benchmarks.

⁴An imperative graph is a mutable data structure where modifications are performed in-place, while a persistent graph is an immutable data structure; see for instance Okasaki's book [14] for more details about persistent data structures.

⁵The Baire library seems to be no longer available from Internet.

1.2 OVERVIEW OF OCAML MODULE SYSTEM

This section quickly describes the OCAML module system. Any reader familiar with OCAML can safely skip this section.

The module system of OCAML is a language by itself, on top of the core OCAML language, which only fulfills software engineering purposes: separate compilation, names space structuring, encapsulation and genericity. This language appears to be independent of the core language [12] and may be unfolded statically. It is a strongly typed higher-order functional language. Its terms are called *modules* or *structures*. They are the basic blocks in OCAML programs, that package together types, values, exceptions and sub-modules.

1.2.1 Structures

Modules are introduced using the `struct...end` construct and the (optional) `module` binding is used to give them a name. Outside a module, its components can be referred to using the *dot notation*: `M.c` denotes the component `c` defined in the module `M`.

For instance, a module packaging together a type for a graph data structure and some basic operations can be implemented in the following way:

```
module Graph = struct
  type label = int
  type t = (int × label) list array
  let create n = Array.create n []
  let add_edge g v1 v2 l = g.(v1) ← (v2,l)::g.(v1)
  let iter_succ g f v = List.iter f g.(v)
end
```

The type `Graph.t` defines a naive graph data structure using adjacency lists with edges labeled with integers: a graph is an array (indexed by integers representing vertices) whose elements are lists of pairs of integers and labels (declared as an alias for the type `int`).

1.2.2 Signatures

The type of a module is called a *signature* or an *interface*⁶ and can be used to hide some components or the definition of a type (then called an *abstract data type*). Signatures are defined using the `sig...end` construct and the (optional) `module type` binding is used to give them a name. Constants and functions are declared via the keyword `val` and types via the keyword `type`.

For instance, a possible signature for the `Graph` module above, that hides the graph representation and the type of labels, could be the following:

```
module type GRAPH = sig
```

⁶with the same meaning as in MODULA but not as in JAVA

```

type label
type t
val create : int → t
val add_edge : t → int → int → label → unit
val iter_succ :
  t → (int × label → unit) → int → unit
end

```

Restricting a structure by a signature results in another view of the structure. This is done as follows:

```

module G' = (G : GRAPH)

```

Since interfaces and structures are clearly separated, it is possible to have several implementations for the same interface. Conversely, a structure may have several signatures (hiding and restricting more or less components).

1.2.3 Functors

The functions of the module system are called *functors* and allow us to define modules parameterized by other modules. Then they can be applied one or several times to particular modules with the expected signature. The benefits of functors in software engineering are appreciated as soon as one has to parameterize a *set* of types and functions by another *set* of types and functions, in a sound way⁷. For instance, to implement Dijkstra's shortest path algorithm for any graph implementation where edges are labeled with integers, one can write a functor looking like:

```

module type S = sig
  type label
  type t
  val iter_succ :
    t → (int × label → unit) → int → unit
end

module Dijkstra (G : S with type label = int) =
  struct
    let dijkstra g v1 v2 = (* ... *)
  end

```

The `with type` annotation is used here to unify the abstract type `label` from the signature `S` with the type `int`. One may also notice that the signature `S` required for the functor's argument only contains what is necessary to implement the algorithm. However, we can apply the functor to any module whose signature contains *at least* `S` i.e. is a *subtype* of `S`.

⁷See for instance Norman Ramsey's *ML Module Mania* [15] as an example of a massive use of ML functors.

```

module G = Imperative.Graph.Abstract
          (struct type t = int × int end)

let g = G.create ()

let () =
  ... add vertices to g with G.add_vertex,
      edges with G.add_edge and initial
      constraints (20 lines of code) ...

module C = Coloring(G)

let () = C.coloring g 9

```

FIGURE 1.2. A Sudoku solver using OCAMLGRAPH

Functors are also first-class values, i.e they can be passed as arguments to other functors. Finally, it is possible to aggregate signatures or modules using the `include` construct which can be naively seen as a textual inclusion.

1.3 MOTIVATING EXAMPLE

To illustrate the use of OCAMLGRAPH, we consider a Sudoku solver based on graph coloring. The idea is to represent the Sudoku grid as an undirected graph with 9×9 vertices, each vertex being connected to all other vertices on the same row, column and 3×3 group. Solving the Sudoku is equivalent to 9-coloring this graph. Figure 1.2 displays the sketch of a solution to this problem using OCAMLGRAPH⁸. There are four steps in this code:

1. We choose a graph data structure for our Sudoku solver: it is an imperative undirected graph with vertices labeled with pairs of integers (the cells coordinates) and unlabeled edges. In this structure, vertices are also equipped with integer marks, that are used to store the assigned colors.
2. We create the Sudoku grid and fill it with the initial constraints.
3. We obtain a coloring algorithm for our graph data structure.
4. We solve the Sudoku problem by 9-coloring the graph.

This code is almost as efficient as a hand-coded Sudoku solver: on the average, a Sudoku puzzle is solved in 0.2 seconds (on a Pentium IV 2.4 GHz). The remainder of this paper goes into more details about the code above.

⁸Full source code for the Sudoku example is given in appendix A.

1.4 SIGNATURES AND GRAPH DATA STRUCTURES

Managing many variants of graph data structures without proper parameterization results into unmanageable code. Here we show how we factorized the OCAML-GRAPH implementation to avoid such pitfall. Section 1.4.1 describes the common sub-signatures shared by all graphs. Section 1.4.2 details their various implementations.

1.4.1 Sharing Signatures for All Graphs

All graph data structures share a common sub-signature G for observers. Two other signatures distinguish the modifiers for persistent and imperative graphs, respectively.

The common signature G includes an abstract type t for the graph datatype and two modules V and E for vertices and edges respectively. The signature for E always includes a type `label` which is instantiated by the singleton type `unit` for unlabeled graphs. Modules V and E both implement the standard comparison and hashing functions so that graph algorithms may easily construct data structures containing vertices and edges. G also includes usual observers such as functions to iterate over vertices and edges, which are massively used in graph algorithms. The common signature looks like:

```
module type G = sig
  type t
  module V : sig type t ... end
  module E : sig
    type t
    type label
    val label : t → label
    ...
  end
  val iter_vertex : (V.t → unit) → t → unit
  val iter_succ : (V.t → unit) → t → V.t → unit
  ...
end
```

We distinguish the signature P for persistent graphs from the signature I for imperative graphs, since the modifiers do not have the same type in both:

```
module type P = sig
  include G
  val empty : t
  val add_vertex : t → V.t → t
  val add_edge : t → V.t → V.t → t
  ...
end
```

```

module type I = sig
  include G
  val create : unit → t
  val add_vertex : t → V.t → unit
  val add_edge : t → V.t → V.t → unit
  ...
end

```

1.4.2 19 Graph Data Structures in 1000 Lines of Code

OCAMLGRAPH provides 19 graph data structures, which include all the possible combinations of the following 4 criteria:

- *directed* or *undirected* graph;
- *labeled* or *unlabeled* edges;
- *persistent* or *imperative* data structure;
- *concrete* or *abstract* type for vertices.

The last point requires some explanations. Vertices are always labeled internally with the value provided by the user. Accessing this value depends on the choice of concrete or abstract vertices. Concrete vertices allow unrestricted access to their value. Abstract vertices hide their value inside an abstract data type. The former allows a more immediate use of the data structure and the latter a more efficient implementation. In particular, imperative graphs with abstract vertices can be equipped with integer mutable marks, which are used in our Sudoku solver.

A functor is provided for each data structure. It is parameterized by user types for vertex labels and possibly edge labels. These functors⁹ are displayed in Figure 1.3 as square boxes mapping signatures of input modules (incoming plain edges) to the signature of the graph module (outgoing plain edges). Of these, 8 functors exist in both directed and undirected versions. Input signatures `ANY_TYPE`, `ORDERED_TYPE_DFT` and `COMPARABLE` define the user types for vertices and edges labels. For instance, functor `AbstractLabelled` from `Imperative.Graph` takes as arguments two modules of signatures `ANY_TYPE` and `ORDERED_TYPE_DFT` respectively and produces a module implementing signature `IM`. This signature extends signature `I` with mutable marks, as indicated by the dashed edge from `IM` to `I`. Three other implementations complete the set of graph data structures, namely `ConcreteBidirectional` for graphs with an efficient access to predecessors, and `Matrix.(Graph, Digraph)` for graphs implemented as adjacency matrices. For efficiency reasons, these three implementations do not offer the same combination of criteria as the previous ones.

Several functors are used internally to avoid code duplication among the functors presented in Figure 1.3. For instance, a functor adds labels to unlabeled graphs; another one encapsulates concrete vertices into an abstract data type; etc.

⁹The signatures of all these functors are available at <http://ocamlgraph.lri.fr/doc/>.

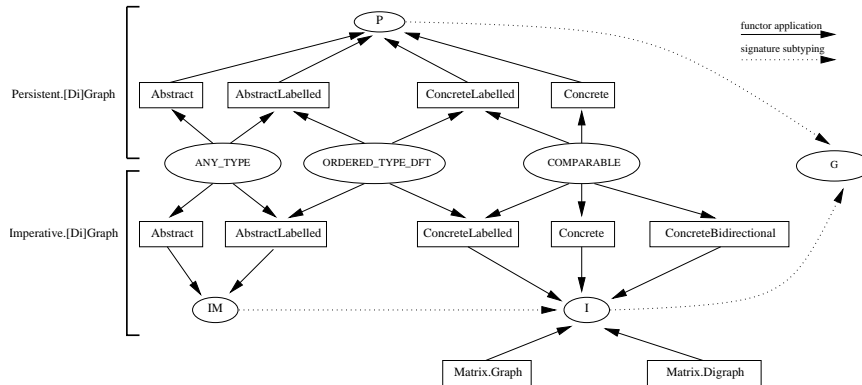


FIGURE 1.3. OCAMLGRAPH data structures components

Putting it all together, the code size for the 19 graph data structures is about 1000 lines. This is clearly small enough to be easily maintained. In Section 1.6 we will show that this code is also quite efficient.

The graph data structure for our Sudoku solver is simply an imperative undirected graph with abstract vertices labeled with pairs of integers and unlabeled edges. It is obtained as:

```
module G = Imperative.Graph.Abstract
      (struct type t = int × int end)
```

1.5 GENERIC ALGORITHMS

This section introduces the second use of functors in OCAMLGRAPH: generic programming of graph algorithms.

1.5.1 Decoupling Algorithms and Graph Data Structures

As demonstrated in Section 1.4, our library provides many graph data structures. It makes it necessary to factorize the code for graph algorithms that operate on these structures. Again, functors provide a nice encoding of generic algorithms.

The basic idea when coding an algorithm is to focus only on the required operations that this algorithm imposes on the graph data structure. Then this algorithm can be expressed naturally as a functor parameterized by these operations. These operations usually form a subset of the operations provided by OCAMLGRAPH graph data structures. In a few cases, the algorithm requires specific operations that are independent of the graph data structure. In such a case, the specific operations are provided as an additional functor parameter.

Such a “functorization” of algorithms has two benefits: first, it allows to add quickly new algorithms to the library, without duplicating code for all data structures; secondly, it allows the user to apply an existing algorithm on his own graph

data structure. Note that on the latter case the user is responsible for fulfilling the requirements over the functor parameters (which are available from OCAML-GRAPH documentation).

1.5.2 Example: Depth-First Traversal

We illustrate the generic programming of graph algorithms on the particular example of depth-first prefix traversal (DFS). To implement DFS, we need to iterate over the graph vertices and over the edges leaving a given vertex. If we do not assume any kind of marks on vertices, we also need to build a data structure to store the visited nodes. We choose a hash table for this purpose and thus we require a hash function and an equality over vertices. Thus the minimal input signature for the DFS functor is as follows:

```
module type G = sig
  type t
  module V : sig
    type t
    val hash : t → int
    val equal : t → t → bool
  end
  val iter_vertex : (V.t → unit) → t → unit
  val iter_succ : (V.t → unit) → t → V.t → unit
end
```

The DFS algorithm is then implemented as a functor with an argument of signature `G`. The result of functor application is a module providing a single function `dfs` to traverse a given graph while applying a given function on all visited vertices:

```
module Dfs(G : G) :
  sig val dfs : (G.V.t → unit) → G.t → unit end
```

To implement this functor, we first instantiate OCAML's generic hash tables on graph vertices:

```
module Dfs(G : G) = struct
  module H = Hashtbl.Make(G.V)
```

Then we can implement the traversal. The following code uses a hash table `h` to store the vertices already visited and an explicit stack `stack` to store the vertices to be visited (to avoid the possible stack overflow of a recursive implementation). Function `G.iter_vertex` is used to start a DFS on every vertex. The DFS itself is performed in function `loop` using `G.iter_succ`:

```
let dfs f g =
  let h = H.create 65537 in
  let stack = Stack.create () in
```

```

let push v =
  if not (H.mem h v) then
    begin H.add h v (); Stack.push v stack end
in
let loop () =
  while not (Stack.is_empty stack) do
    let v = Stack.pop stack in
      f v;
      G.iter_succ push g v
  done
in
G.iter_vertex (fun v → push v; loop ()) g
end

```

Beside this simple algorithm, OCAMLGRAPH provides other kinds of traversals (breadth-first, postfix, etc.) and more efficient implementations when the graph data structure contains mutable marks on vertices.

1.5.3 Example: Graph Coloring

As a second example, we present a graph coloring algorithm used in our Sudoku solver. For the purpose of our algorithm, we require the presence of `get` and `set` operations on integer marks associated to vertices. We use these marks to store the color assigned to each vertex. We also need iterators over vertices and successors. Thus the minimal signature for a graph data structure used in our graph coloring algorithm is the following:

```

module type GM = sig
  type t
  module V : sig type t ... end
  module Mark : sig
    val get : V.t → int
    val set : V.t → int → unit
  end
  val iter_vertex : (V.t → unit) → t → unit
  val iter_succ : (V.t → unit) → t → V.t → unit
end

```

OCAMLGRAPH already provides implementations for such a signature. This is the case for the graph data structure used in our Sudoku solver. Then the graph coloring algorithm is implemented as the following functor:

```

module Coloring(G : GM) : sig
  val coloring : G.t → int → unit
end

```

It provides a single function `coloring` which colors a given graph with a given number of colors. Some marks may contain initial constraints. The implementa-

tion of this coloring algorithm is given in appendix B. To complete our Sudoku solver, we simply need to apply the above functor on our graph module `G`:

```
module C = Coloring(G)
```

If `g` contains the Sudoku graph, and assuming that the initial constraints are set in `g` marks, solving the Sudoku amounts to 9-coloring graph `g`:

```
C.coloring g 9
```

1.5.4 Building Graphs

In Section 1.4.1, we have shown that persistent and imperative graphs have creation functions with different signatures. However, as we have written algorithms in a generic way, we may want to build graphs in a generic way, that is independently of the underlying data structure. For instance, we may want to implement graph operations (such as union, transitive closure, etc.) or to build some classic graphs (such as the full graph with n vertices, the de Bruijn graph of order n , etc.) or even random graphs. In all these cases, the persistent or imperative nature of the graph is not really significant but the signature difference disallows genericity.

To solve this issue, we introduce a module `Builder`. It defines a common interface for graphs building:

```
module type S = sig
  module G : Sig.G
  val empty : unit → G.t
  val copy : G.t → G.t
  val add_vertex : G.t → G.V.t → G.t
  val add_edge : G.t → G.V.t → G.V.t → G.t
  val add_edge_e : G.t → G.E.t → G.t
end
```

It is immediate to realize such a signature for persistent or imperative graphs:

```
module P(G : Sig.P) : S with module G = G
module I(G : Sig.I) : S with module G = G
```

It is important to notice that for imperative graphs the values returned by the functions `add_vertex`, `add_edge` and `add_edge_e` are meaningless.

Therefore, it is easy to write a generic algorithm that builds graphs. First we write a generic version as a functor taking a module of signature `Builder.S` as argument:

```
module Make(B : Builder.S) = struct ... end
```

and then we can trivially provide two variants of this functor for both persistent and imperative graphs, with the following two lines:

```
module P(G : Sig.P) = Make(Builder.P(G))
module I(G : Sig.I) = Make(Builder.I(G))
```

Thus the use of the module `Builder` is entirely hidden from the user point of view.

1.6 BENCHMARKS

Surprisingly, we could not find any standard benchmark for graph libraries. In order to give an idea of OCAMLGRAPH efficiency, we present here the results of a little benchmark of our own. We test four different data structures for undirected graphs with unlabeled edges, that are either persistent (P) or imperative (I) and with either abstract (A) or concrete (C) vertices. In the following, these are referred to as PA, PC, IA and IC, respectively. All tests were performed on a Pentium 4 2.4 GHz.

We first test the efficiency of graph creation and mutation. For that purpose, we build cliques of V vertices (and thus $E = V(V + 1)/2$ edges since we include self loops). Then we repeatedly delete all edges and vertices in these graphs. Figure 1.4 displays the creation and deletion timings in seconds up to $V = 1000$ (that is half a million edges). The speed of creation observed is roughly 100,000 edges per second for imperative graphs. The creation of persistent graphs is slower but within a constant factor (less than 2). Deletion is twice as fast as creation. Regarding memory consumption, all four data structures use approximately 5 machine words (typically 20 bytes) per edge.

Our second benchmark consists in generating graphs corresponding to 2D mazes and traversing them using depth-first and breadth-first traversals. Given an integer N , we build a graph with $V = N^2$ vertices and $E = V - 1$ edges. Figure 1.5 displays the timings in seconds for various values of N up to 600 (i.e. 360,000 vertices). The observed speed is between 500,000 and 1 million traversed edges per second.

We also tested the adjacency matrix-based data structure. Creation and deletion are much faster in that case, and the data structure for a dense graph is usually much more compact (it is implemented using bit vectors). However, the use of this particular implementation is limited to unlabeled imperative graphs with integer vertices. The above benchmarks, on the contrary, do not depend on the nature of vertices and edges types. Thus they are much more representative of OCAMLGRAPH average performances.

1.7 CONCLUSION

We presented OCAMLGRAPH, a highly generic graph library for OCAML providing several graph data structures and graph algorithms. Algorithms are written independently from graph data structures, which allows combining user data structure (resp. algorithm) with OCAMLGRAPH algorithm (resp. data structure). To our knowledge, there is no library for any applicative language as generic as OCAMLGRAPH. This genericity is obtained using OCAML module system and especially its functors which allow sharing large pieces of code and provide clear separation between data structures and algorithms. The same level of genericity could probably be achieved using Haskell's multi-parameter type classes [8, 16, 11]. Regarding imperative languages, graph libraries are rarely as generic and never provide as many different data structures.

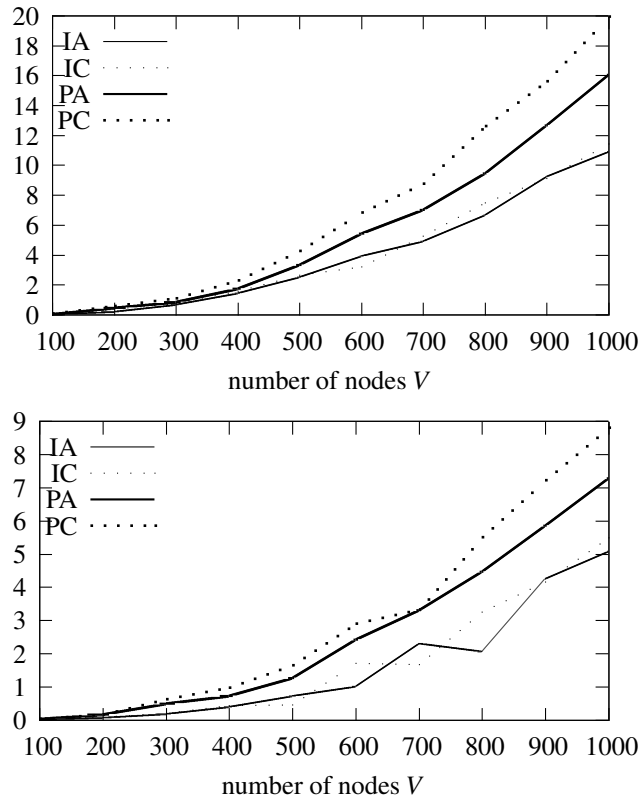


FIGURE 1.4. Benchmarking creation (top) and deletion (bottom)

Since its first release (Feb. 2004), the number of OCAMLGRAPH users has been increasing steadily and several of them contributed code to the library. Some of them provided new graph data structures (*e.g.* `ConcreteBidirectional`) and others new algorithms (*e.g.* minimal separators). It clearly shows the benefits of a generic library where data structures and algorithms are separated.

ACKNOWLEDGEMENTS

We would like to thank the anonymous reviewers for their helpful comments. We already improved the implementation of OCAMLGRAPH following one of the reviewer suggestions (generalizing mutable marks from integer to arbitrary types).

REFERENCES

- [1] Baire. <http://www.edite-de-paris.com.fr/~fpons/Caml/Baire/>.




FIGURE 1.5. Benchmarking DFS (top) and BFS (bottom)

- [2] BGL - The Boost Graph Library. <http://www.boost.org/libs/graph/doc/>.
- [3] FGL - A Functional Graph Library. <http://web.engr.oregonstate.edu/~erwig/fgl/>.
- [4] The Data Structures Library in Java. <http://www.cs.brown.edu/cgc/jdsl/>.
- [5] The Graph Template Library. <http://infosun.fmi.uni-passau.de/GTL/>.
- [6] The MLRISC System. <http://cs1.cs.nyu.edu/leunga/www/MLRISC/Doc/html/INTRO.html>.
- [7] The Objective Caml language. <http://caml.inria.fr/>.
- [8] D. Dreyer, R. Harper, and M. M. T. Chakravarty. Modular type classes. In *POPL*, 2007.
- [9] Martin Erwig. Inductive graphs and functional graph algorithms. *Journal of Functional Programming*, 11(5):467–492, 2001.

- [10] Jean-Christophe Filliâtre. Backtracking Iterators. Research Report 1428, LRI, Université Paris-Sud, January 2006. <http://www.lri.fr/~filliatr/ftp/publis/enum-rr.ps.gz>.
- [11] Oleg Kiselyov. Applicative translucent functors in Haskell, 2004. At <http://www.haskell.org/pipermail/haskell/2004-August/014463.html>.
- [12] Xavier Leroy. A modular module system. *Journal of Functional Programming*, 10(3):269–303, 2000.
- [13] Kurt Mehlhorn and Stefan Nher. Leda: a platform for combinatorial and geometric computing. *Commun. ACM*, 38(1):96–102, 1995.
- [14] Chris Okasaki. *Purely Functional Data Structures*. Cambridge University Press, 1998.
- [15] Norman Ramsey. ML Module Mania: A Type-Safe Separately Compiled, Extensible Interpreter. In *ACM SIGPLAN Workshop on ML*, 2005.
- [16] Stefan Wehr. *ML Modules and Haskell Type Classes: A Constructive Comparison*, November 2005. Submitted for publication and available at <http://www.stefanwehr.de/diplom>.

A SIMPLE SUDOKU SOLVER USING OCAMLGRAPH

Below is the full listing for a Sudoku solver using OCAMLGRAPH, as described in this paper. This program reads the Sudoku problem on standard input and prints the solution on standard output.

```
open Graph

(* We use undirected graphs with nodes containing
   a pair of integers (the cell coordinates in
   0..8 x 0..8). *)
module G = Imperative.Graph.Abstract
        (struct type t = int * int end)

(* The Sudoku grid = a graph with 9x9 nodes *)
let g = G.create ()

(* We create the 9x9 nodes, add them to the graph
   and keep them in a matrix for later access *)
let nodes =
  let new_node i j =
    let v = G.V.create (i, j) in G.add_vertex g v; v
  in
  Array.init 9 (fun i -> Array.init 9 (new_node i))

let node i j = nodes.(i).(j)

(* We add the edges: two nodes are connected whenever
   they can't have the same value *)
let () =
```

```

for i = 0 to 8 do for j = 0 to 8 do
  for k = 0 to 8 do
    if k <> i then G.add_edge g (node i j) (node k j);
    if k <> j then G.add_edge g (node i j) (node i k);
  done;
  let gi = 3 * (i / 3) and gj = 3 * (j / 3) in
  for di = 0 to 2 do for dj = 0 to 2 do
    let i' = gi + di and j' = gj + dj in
    if i' <> i || j' <> j then
      G.add_edge g (node i j) (node i' j')
  done done
done done

(* We read the initial constraints from standard input *)
let () =
  for i = 0 to 8 do
    let s = read_line () in
    for j = 0 to 8 do match s.[j] with
      | '1'..'9' as ch ->
        G.Mark.set (node i j) (Char.code ch - Char.code '0')
      | _ -> ()
    done
  done

(* We solve the Sudoku by 9-coloring the graph g *)
module C = Coloring.Mark(G)
let () = C.coloring g 9

(* We display the solution *)
let () =
  for i = 0 to 8 do
    for j = 0 to 8 do
      Format.printf "%d" (G.Mark.get (node i j))
    done;
    Format.printf "\n";
  done;
  Format.printf "@?"

```

B GRAPH COLORING IMPLEMENTATION

Below is the code of the graph coloring functor introduced in Section 1.5.3 and used to write the Sudoku solver. The code uses a simple backtracking algorithm which performs a breadth-first traversal of the graph and successively tries each color for each visited vertex. To be able to backtrack during the traversal we use persistent cursors [10] provided by another OCAMLGRAPH functor, namely `Traverse.Bfs`. A persistent cursor is created with function `start`, the visited element is obtained with function `get` and the cursor is moved to the next element with function `step`. The latter returns a new cursor, contrary to usual cursors which are modified in-place, thus allowing backtracking.


```

module Coloring(G : GM) = struct
  module Bfs = Traverse.Bfs(G)

  exception NoColoring

  let coloring g k =
    (* assign color i to vertex v, if possible,
       and raise NoColoring otherwise *)
    let try_color v i =
      G.iter_succ
        (fun w ->
          if G.Mark.get w = i then raise NoColoring)
      g v;
      G.Mark.set v i
    in
    (* traversal of g using persistent cursor iter *)
    let rec iterate iter =
      let v = Bfs.get iter in
      for i = 1 to k do
        try try_color v i; iterate (Bfs.step iter);
          assert false
        with NoColoring -> ()
      done;
      G.Mark.set v 0; raise NoColoring
    in
    try iterate (Bfs.start g); assert false
    with Exit -> ()
  end
end

```

Note that the actual implementation in OCAMLGRAPH is slightly more complex since it uses Kempe's linear-time simplification (vertices of degree less than k are repeatedly removed and pushed on a stack, then the backtracking algorithm above is performed and finally vertices initially removed are popped from the stack and colored).