The Boost Graph Library

User Guide and Reference Manual

Jeremy G. Siek Lie-Quan Lee Andrew Lumsdaine

Foreword by Alexander Stepanov

C++ In-Depth Series • Bjarne Stroustrup

Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this book, and Addison-Wesley was aware of a trademark claim, the designations have been printed with initial capital letters or in all capitals.

The authors and publisher have taken care in the preparation of this book, but make no expressed or implied warranty of any kind and assume no responsibility for errors or omissions. No liability is assumed for incidental or consequential damages in connection with or arising out of the use of the information or programs contained herein.

The publisher offers discounts on this book when ordered in quantity for special sales. For more information, please contact:

Pearson Education Corporate Sales Division One Lake Street Upper Saddle River, NJ 07458 (800) 382-3419 corpsales@pearsontechgroup.com

Visit AW on the Web: www.aw.com/cseng/

Library of Congress Cataloging-in-Publication Data

Siek, Jeremy G.
The Boost graph library : user guide and reference manual/ Jeremy G. Siek, Lie-Quan Lee, Andrew Lumsdaine
p. cm.
Includes bibliographical references and index.
ISBN 0-201-72914-8 (alk. paper)
1. C++ (Computer language). 2. Graph theory. I. Lee, Lie-Quan.
II. Lumsdaine, Andrew. III. Title.

006.6-dc21

2001053553

Copyright © 2002 Pearson Education, Inc.

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form, or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior consent of the publisher. Printed in the United States of America. Published simultaneously in Canada.

ISBN 0-201-72914-8 Text printed on recycled paper 1 2 3 4 5 6 7 8 9 10—MA—0504030201 First printing, December 2001

Contents

Foreword				xiii xvii	
Preface					
Ι	Use	r Guid	e	1	
1	Intr	oductio	n	3	
	1.1	Some	Graph Terminology	3	
	1.2	Graph	Concepts	5	
		1.2.1	Vertex and Edge Descriptors	5	
		1.2.2	Property Maps	6	
		1.2.3	Graph Traversal	7	
		1.2.4	Graph Construction and Modification	9	
		1.2.5	Algorithm Visitors	10	
	1.3	Graph	Classes and Adaptors	11	
		1.3.1	Graph Classes	11	
		1.3.2	Graph Adaptors	13	
	1.4	Generi	c Graph Algorithms	13	
		1.4.1	The Topological Sort Generic Algorithm	14	
		1.4.2	The Depth-First Search Generic Algorithm	18	
2	Gen	eric Pro	ogramming in C++	19	
	2.1	Introdu	action	19	
		2.1.1	Polymorphism in Object-Oriented Programming	20	
		2.1.2	Polymorphism in Generic Programming	21	
		2.1.3	Comparison of GP and OOP	22	
	2.2	Generi	c Programming and the STL	25	
	2.3	Conce	pts and Models	27	
		2.3.1	Sets of Requirements	28	
		2.3.2	Example: InputIterator	28	

	2.4	Associated Types and Traits Classes	30				
		2.4.1 Associated Types Needed in Function Template	30				
		2.4.2 Typedefs Nested in Classes	30				
		2.4.3 Definition of a Traits Class	31				
		2.4.4 Partial Specialization	32				
		2.4.5 Tag Dispatching	33				
	2.5	Concept Checking	34				
		2.5.1 Concept-Checking Classes	35				
		2.5.2 Concept Archetypes	36				
	2.6	The Boost Namespace	37				
		2.6.1 Classes	37				
		2.6.2 Koenig Lookup	38				
	2.7	Named Function Parameters	39				
2			41				
3	A B(JL IUTOFIAI	41				
	3.1	Graph Satur	41 42				
	3.2	Compilation Order	42				
	5.5	3 3 1 Topological Sort via DES	44 11				
		3.3.2 Marking Vartices Using External Properties	44				
		3.3.3 Accessing Adjacent Vertices	40				
		3.3.4 Traversing All the Vertices	40				
	31	Cyclic Dependencies	47				
	3.5	Toward a Generic DES: Visitors	-10 40				
	3.6	Granh Setun: Internal Properties					
	37	Compilation Time					
	3.8	A Generic Topological Sort and DES	55				
	3.9	Parallel Compilation Time	57				
	3.10	Summary	59				
	0.110						
4	Basi	c Graph Algorithms	61				
	4.1	Breadth-First Search	61				
		4.1.1 Definitions	61				
		4.1.2 Six Degrees of Kevin Bacon	62				
	4.2	Depth-First Search	67				
		4.2.1 Definitions	67				
		4.2.2 Finding Loops in Program-Control-Flow Graphs	69				
5	Shor	test-Paths Problems	75				
-	5.1	Definitions	75				
	5.2	Internet Routing	76				
	_	σ	2				

	5.3 5.4	Bellman–Ford and Distance Vector Routing77Dijkstra and Link-State Routing81
6	Mini 6.1 6.2 6.3 6.4	mum-Spanning-Tree Problem89Definitions89Telephone Network Planning89Kruskal's Algorithm91Prim's Algorithm94
7	Coni 7.1 7.2 7.3	Definitions97Definitions97Connected Components and Internet Connectivity98Strongly Connected Components and Web Page Links102
8	Max 8.1 8.2	imum Flow105Definitions105Edge Connectivity106
9	Impl 9.1 9.2 9.3	icit Graphs: A Knight's Tour113Knight's Jumps as a Graph114Backtracking Graph Search116Warnsdorff's Heuristic117
10	Inter 10.1 10.2 10.3	facing with Other Graph Libraries119Using BGL Topological Sort with a LEDA Graph120Using BGL Topological Sort with a SGB Graph122Implementing Graph Adaptors123
11	Perf 11.1 11.2	Ormance Guidelines127Graph Class Comparisons12711.1.1 The Results and Discussion128Conclusion132
II	Ref	Cerence Manual 135
12	BGL 12.1	Concepts 137 Graph Traversal Concepts 137 12.1.1 Undirected Graphs 138 12.1.2 Graph 142 12.1.3 IncidenceGraph 143 12.1.4 BidirectionalGraph 145

		12.1.5	AdjacencyGraph	6
		12.1.6	VertexListGraph	7
		12.1.7	EdgeListGraph	8
		12.1.8	AdjacencyMatrix	9
	12.2	Graph	Modification Concepts	0
		12.2.1	VertexMutableGraph	2
		12.2.2	EdgeMutableGraph	2
		12.2.3	MutableIncidenceGraph	4
		12.2.4	MutableBidirectionalGraph	4
		12.2.5	MutableEdgeListGraph	5
		12.2.6	PropertyGraph	5
		12.2.7	VertexMutablePropertyGraph	6
		12.2.8	EdgeMutablePropertyGraph	7
	12.3	Visitor	Concepts	8
		12.3.1	BFSVisitor	8
		12.3.2	DFSVisitor	0
		12.3.3	DijkstraVisitor	1
		12.3.4	BellmanFordVisitor	2
13	BGL	Algori	thms 16	3
			10	<i>U</i>
	13.1	Overvi	ew	3
	13.1 13.2	Overvi Basic A	ew	3 5
	13.1 13.2	Overvi Basic A 13.2.1	ew 16. Algorithms 16. breadth_first_search 16.	3 5 5
	13.1 13.2	Overvi Basic A 13.2.1 13.2.2	ew 16 Algorithms 16 breadth_first_search 16 breadth_first_visit 16	3 5 5 9
	13.1 13.2	Overvi Basic A 13.2.1 13.2.2 13.2.3	ew 16. Algorithms 16. breadth_first_search 16. breadth_first_visit 16. breadth_first_visit 16. 16. 16. breadth_first_search 16. 16. 16. 16. 16. 16. 16. 16. 16. 16. 16. 16. 16. 16. 16. 16. 170.	2 3 5 5 9 0
	13.1 13.2	Overvi Basic A 13.2.1 13.2.2 13.2.3 13.2.4	ew16.Algorithms16.breadth_first_search16.breadth_first_visit16.breadth_first_visit16.depth_first_search17.depth_first_visit17.17.17.17.17.17.17.	3 5 5 9 0 5
	13.1 13.2	Overvi Basic A 13.2.1 13.2.2 13.2.3 13.2.4 13.2.5	ew 16 Algorithms 16 breadth_first_search 16 breadth_first_visit 16 breadth_first_visit 16 depth_first_search 16 depth_first_search 17 depth_first_visit 17 topological_sort 17	3 5 5 9 0 5 6
	13.1 13.2 13.3	Overvi Basic A 13.2.1 13.2.2 13.2.3 13.2.4 13.2.5 Shortes	ew 16. Algorithms 16. breadth_first_search 16. breadth_first_visit 16. breadth_first_visit 16. depth_first_visit 16. depth_first_search 17. depth_first_visit 17. topological_sort 17. st-Path Algorithms 17.	0 3 5 5 9 0 5 6 7
	13.1 13.2 13.3	Overvi Basic A 13.2.1 13.2.2 13.2.3 13.2.4 13.2.5 Shortes 13.3.1	ew 16 Algorithms 16 breadth_first_search 16 breadth_first_visit 16 breadth_first_visit 16 depth_first_search 16 depth_first_visit 17 depth_first_visit 17 topological_sort 17 st-Path Algorithms 17 dijkstra_shortest_paths 17	3 5 5 9 0 5 6 7 7
	13.1 13.2 13.3	Overvii Basic A 13.2.1 13.2.2 13.2.3 13.2.4 13.2.5 Shortes 13.3.1 13.3.2	ew 16 Algorithms 16 breadth_first_search 16 breadth_first_visit 16 breadth_first_visit 16 depth_first_visit 16 depth_first_search 16 depth_first_visit 17 depth_first_visit 17 topological_sort 17 st-Path Algorithms 17 dijkstra_shortest_paths 17 bellman_ford_shortest_paths 18	2 3 5 5 9 0 5 6 7 7 2
	13.1 13.2 13.3	Overvi Basic A 13.2.1 13.2.2 13.2.3 13.2.4 13.2.5 Shortes 13.3.1 13.3.2 13.3.3	ew 16. Algorithms 16. breadth_first_search 16. breadth_first_visit 16. breadth_first_visit 16. breadth_first_visit 16. depth_first_search 16. depth_first_visit 17. depth_first_visit 17. topological_sort 17. st-Path Algorithms 17. dijkstra_shortest_paths 17. bellman_ford_shortest_paths 18. johnson_all_pairs_shortest_paths 18.	3 5 5 9 0 5 6 7 7 2 6
	13.113.213.313.4	Overvii Basic A 13.2.1 13.2.2 13.2.3 13.2.4 13.2.5 Shortes 13.3.1 13.3.2 13.3.3 Minimu	ew 16 Algorithms 16 breadth_first_search 16 breadth_first_visit 16 breadth_first_visit 16 depth_first_search 16 depth_first_visit 17 depth_first_visit 17 topological_sort 17 st-Path Algorithms 17 dijkstra_shortest_paths 17 bellman_ford_shortest_paths 18 johnson_all_pairs_shortest_paths 18 um-Spanning-Tree Algorithms 18	2355905677269
	13.1 13.2 13.3 13.4	Overvii Basic A 13.2.1 13.2.2 13.2.3 13.2.4 13.2.5 Shortes 13.3.1 13.3.2 13.3.3 Minimi 13.4.1	ew 16 Algorithms 16 breadth_first_search 16 breadth_first_visit 16 breadth_first_visit 16 depth_first_search 16 depth_first_visit 16 depth_first_visit 17 depth_first_visit 17 topological_sort 17 st-Path Algorithms 17 dijkstra_shortest_paths 17 bellman_ford_shortest_paths 18 johnson_all_pairs_shortest_paths 18 kruskal_minimum_spanning_tree 18	3559056772699
	13.1 13.2 13.3 13.4	Overvi Basic A 13.2.1 13.2.2 13.2.3 13.2.4 13.2.5 Shortes 13.3.1 13.3.2 13.3.3 Minimu 13.4.1 13.4.2	ew 16 Algorithms 16 breadth_first_search 16 breadth_first_visit 16 breadth_first_visit 16 depth_first_search 16 depth_first_visit 17 depth_first_visit 17 topological_sort 17 st-Path Algorithms 17 dijkstra_shortest_paths 17 bellman_ford_shortest_paths 18 johnson_all_pairs_shortest_paths 18 wm-Spanning-Tree Algorithms 18 prim_minimum_spanning_tree 18 prim_minimum_spanning_tree 19	2 3 5 5 9 0 5 6 7 7 2 6 9 9 2
	13.1 13.2 13.3 13.4 13.5	Overvii Basic A 13.2.1 13.2.2 13.2.3 13.2.4 13.2.5 Shortes 13.3.1 13.3.2 13.3.3 Minimu 13.4.1 13.4.2 Static O	ew 16 Algorithms 16 breadth_first_search 16 breadth_first_visit 16 breadth_first_visit 16 depth_first_visit 16 depth_first_visit 16 depth_first_visit 16 depth_first_visit 17 depth_first_visit 17 topological_sort 17 st-Path Algorithms 17 dijkstra_shortest_paths 17 bellman_ford_shortest_paths 18 johnson_all_pairs_shortest_paths 18 um-Spanning-Tree Algorithms 18 prim_minimum_spanning_tree 19 Connected Components 19	355905677269925
	13.1 13.2 13.3 13.4 13.5	Overvii Basic A 13.2.1 13.2.2 13.2.3 13.2.4 13.2.5 Shortes 13.3.1 13.3.2 13.3.3 Minimu 13.4.1 13.4.2 Static C 13.5.1	ew 16 Algorithms 16 breadth_first_search 16 breadth_first_visit 16 breadth_first_visit 16 depth_first_search 16 depth_first_visit 16 depth_first_visit 17 depth_first_visit 17 topological_sort 17 st-Path Algorithms 17 dijkstra_shortest_paths 17 bellman_ford_shortest_paths 18 johnson_all_pairs_shortest_paths 18 kruskal_minimum_spanning_tree 19 Connected Components 19 connected_components 19	3559056772699255
	13.1 13.2 13.3 13.4 13.5	Overvi Basic A 13.2.1 13.2.2 13.2.3 13.2.4 13.2.5 Shortes 13.3.1 13.3.2 13.3.3 Minimu 13.4.1 13.4.2 Static C 13.5.1 13.5.2	ew 16. Algorithms 16. breadth_first_search 16. breadth_first_visit 16. breadth_first_visit 16. breadth_first_visit 16. breadth_first_visit 16. breadth_first_visit 16. breadth_first_visit 16. depth_first_visit 17. topological_sort 17. topological_sort 17. topological_sort 17. dijkstra_shortest_paths 17. bellman_ford_shortest_paths 17. bellman_ford_shortest_paths 18. johnson_all_pairs_shortest_paths 18. pinnspanning_Tree 18. prim_minimum_spanning_tree 19. Connected Components 19. strong_components 19. strong_components 19.	0 3 5 5 9 0 5 6 7 7 2 6 9 9 2 5 5 8
	 13.1 13.2 13.3 13.4 13.5 13.6 	Overvii Basic A 13.2.1 13.2.2 13.2.3 13.2.4 13.2.5 Shortes 13.3.1 13.3.2 13.3.3 Minimu 13.4.1 13.4.2 Static C 13.5.1 13.5.2 Increm	ew16.Algorithms16.breadth_first_search16.breadth_first_visit16.depth_first_search16.depth_first_search17.depth_first_visit17.topological_sort17.st-Path Algorithms17.dijkstra_shortest_paths17.bellman_ford_shortest_paths18.johnson_all_pairs_shortest_paths18.prim_minimum_spanning_tree19.Connected Components19.strong_components19.ental Connected Components20.	3 5 5 9 0 5 6 7 7 2 6 9 9 2 5 5 8 1
	 13.1 13.2 13.3 13.4 13.5 13.6 	Overvii Basic A 13.2.1 13.2.2 13.2.3 13.2.4 13.2.5 Shortes 13.3.1 13.3.2 13.3.3 Minimu 13.4.1 13.4.2 Static C 13.5.1 13.5.2 Increm 13.6.1	ew16Algorithms16breadth_first_search16breadth_first_visit16depth_first_visit16depth_first_search17depth_first_visit17topological_sort17dijkstra_shortest_paths17bellman_ford_shortest_paths18johnson_all_pairs_shortest_paths18prim_minimum_spanning_tree19connected Components19strong_components19strong_components20initialize_incremental_components20	3 5 5 9 0 5 6 7 7 2 6 9 9 2 5 5 8 1 3
	13.1 13.2 13.3 13.4 13.5 13.6	Overvi Basic A 13.2.1 13.2.2 13.2.3 13.2.4 13.2.5 Shortes 13.3.1 13.3.2 13.3.3 Minimu 13.4.1 13.4.2 Static C 13.5.1 13.5.2 Increm 13.6.1 13.6.2	ew16Algorithms16breadth_first_search16breadth_first_visit16depth_first_visit17depth_first_visit17topological_sort17st-Path Algorithms17dijkstra_shortest_paths17bellman_ford_shortest_paths18johnson_all_pairs_shortest_paths18prim_minimum_spanning_tree19Connected Components19strong_components19ental Connected Components20initialize_incremental_components20incremental_components20	0 3 5 5 9 0 5 6 7 7 2 6 9 9 2 5 5 8 1 3 3

		13.6.3	same_component
		13.6.4	component_index
	13.7	Maxim	um-Flow Algorithms
		13.7.1	edmunds_karp_max_flow
		13.7.2	push_relabel_max_flow 209
14	BGI	Classe	s 213
	14.1	Graph	Classes
	11	14.1.1	adiacency list 213
		14.1.2	adjacency matrix 235
	14.2	Auxilia	ary Classes
		14.2.1	graph traits
		14.2.2	adiacency list traits
		14.2.3	adjacency matrix traits
		14.2.4	property map
		14.2.5	property
	14.3	Graph	Adaptors
		14.3.1	edge list
		14.3.2	reverse_graph
		14.3.3	$filtered_graph$
		14.3.4	SGB Graph Pointer
		14.3.5	LEDA $GRAPH < V, E > \dots \dots$
		14.3.6	std::vector <edgelist></edgelist>
15	Pron	ertv M	an Library 277
10	15 1	Propert	ty Man Concepts 278
	15.1	15 1 1	BeadablePropertyMap 279
		15.1.1	WritablePropertyMap 280
		15.1.2	BeadWritePropertyMap 281
		15.1.5	I valuePropertyMap 281
	15.2	Propert	tv Map Classes
	15.2	15.2.1	property traits
		15.2.2	iterator property man
		15.2.3	Property Tags
	15.3	Creatin	g Your Own Property Maps
	1010	15.3.1	Property Maps for Stanford GraphBase
		15.3.2	A Property Map Implemented with <i>std::map</i>
16	Anvi	liarv Ca	oncents. Classes, and Functions 280
Ĩ	16.1	Buffer	289 289
	16.2	ColorV	alue 200
	10.4	00101	

16.3	MultiPassInputIterator	291
16.4	Monoid	291
16.5	$mutable_queue$	292
16.6	Disjoint Sets	:93
	16.6.1 <i>disjoint_sets</i>	:93
	16.6.2 find_with_path_halving	:95
	16.6.3 find_with_full_path_compression	:95
16.7	<i>tie</i>	295
16.8	graph_property_iter_range 2	:97
Bibliogr	aphy 2	99
Index	3	603

Foreword

When I first looked at this book, I felt envious. After all, what led me to the discovery of generic programming was the desire to build a library like the Boost Graph Library (BGL). In 1984 I joined the faculty of Polytechnic University in Brooklyn with some vague ideas about building libraries of software components. Well, to tell you the truth that was my secondary interest—my real interest at that time was to construct formal underpinnings of natural language, something like Aristotle's Organon, but more complete and formal. I was probably the only assistant professor in any Electrical Engineering or Computer Science department who meant to obtain tenure through careful study of Aristotle's Categories. Interestingly enough, the design of the Standard Template Library (STL)—in particular the underlying ontology of objects—is based on my realization that the whole-part relation is a fundamental relation that describes the real world and that it is not at all similar to the element-set relation familiar to us from set theory. Real objects do not share parts: my leg is nobody else's leg. STL containers are like that: two containers do not share parts. There are operations like *std::list::splice* that move parts from one container to another; they are similar to organ transplant: my kidney is mine until it is spliced into somebody else.

In any case, I was firmly convinced that software components should be functional in nature and based on John Backus's FP system. The only novel intuition was that functions should be associated with some axioms: for example, the "Russian peasant algorithm" that allows one to compute the *n*th power in $O(\log n)$ steps is defined for any object that has an associative binary operation defined on it. In other words, I believed that algorithms should be associated with what we now call concepts (see §2.3 of this book), but what I called structure types and what type-theorists call *multi-sorted algebras*.

It was my great luck that Polytechnic had a remarkable person on its faculty, Aaron Kershenbaum, who combined deep knowledge of graph algorithms with an unusual desire to implement them. Aaron saw potential in my attempts to decompose programs into simple primitives, and spent a lot of time teaching me graph algorithms and working with me on implementing them. He also showed me that there were some fundamental things that cannot be done functionally without prohibitive change in the complexity. Although it was often possible for me to implement linear time algorithms functionally without changing the asymptotic complexity, it was impossible in practice to implement logarithmic time algorithms without making them linear. In particular, Aaron explained to me why priority queues were so important for many graph algorithms (and he was well qualified to do so: Knuth in his Stanford GraphBase book [22] attributes the discovery of how to apply binary heaps to Prim's and Dijkstra's algorithms to Aaron).

It was a moment of great joy when we were able to produce Prim's and Dijkstra's algorithms as two instances of the same generic—we called it "high-order" then—algorithm. It is quite remarkable how close BGL code is to what we had (see, for example, a footnote to $\S13.4.2$). The following code in Scheme shows how the two algorithms were implemented in terms of the same higher-order algorithm. The only difference is in how distance values are combined: using addition for Dijkstra's and by selecting the second operand for Prim's.

```
(define dijkstra
  (make-scan-based-algorithm-with-mark
    make-heap-with-membership-and-values + < ))</pre>
```

(define prim (make-scan-based-algorithm-with-mark make-heap-with-membership-and-values (lambda (x y) y) <))

It took me a long time—almost 10 years—to find a language in which this style of programming could be *effectively* realized. I finally found C++, which enabled me to produce something that people could use. Moreover, C++ greatly influenced my design by providing a crisp C-based machine model. The features of C++ that enabled STL are templates and overloading.

I often hear people attacking C++ overloading, and, as is true with most good mechanisms, overloading can be misused. But it is an essential mechanism for the development of useful abstractions. If we look at mathematics, it has been greatly driven by overloading. Extensions of a notion of numbers from natural numbers to integers, to rational numbers, to Gaussian integers, to p-adic numbers, etc, are examples of overloading. One can easily guess things without knowing exact definitions. If I see an expression that uses both addition and multiplication, I assume distributivity. If I see less-than and addition, I assume that if a < b then a + c < b + c (I seldom add uncountable cardinals). Overloading allows us to carry knowledge from one type to another.

It is important to understand that one can write generic algorithms just with overloading, without templates: it does, however, require a lot of typing. That is, for every class that satisfies, say, random access iterator requirements, one has to define all the relevant algorithms by hand. It is tedious, but can be done (only signatures would need to be defined: the bodies will be the same). It should be noted that generics in Ada require hand-instantiation and, therefore, are not that helpful, since every algorithm needs to be instantiated by hand. Templates in C++ solve this problem by allowing one to define things once.

There are still things that are needed for generic programming that are not yet representable in C++. Generic algorithms are algorithms that work on objects with similar interfaces. Not identical interfaces as in object-oriented programming, but similar. It is not just the handling of binary methods (see §2.1.3) that causes the problem, it is the fact that interfaces are described in terms of a single type (single-sorted algebra). If we look carefully at things like iterators we observe that they are describable only in terms of multiple types: the iterator type itself, the value type, and the distance type. In other words, we need three types to define the interfaces on one type. And there is no machinery in C++ to do that. The result of this is that we cannot define what iterators are and, therefore, cannot really compile generic algorithms. For example, if we define the reduce algorithm as:

```
template <class InputIterator, class BinaryOperationWithIdentity>
typename iterator_traits<InputIterator>::value_type
reduce(InputIterator first, InputIterator last, BinaryOperationWithIdentity op)
{
   typedef typename iterator_traits<InputIterator>::value_type T;
   if (first == last) return identity_element(op);
   T result = *first;
   while (++first != last) result = op(result, *first);
   return result;
}
```

but instead of: ++first != last we write: ++first < last, no compiler can detect the bug at the point of definition. Though the standard clearly states that *operator* < does not need to be defined for Input Iterators, there is no way for the compiler to know it. Iterator requirements are just words. We are trying to program with concepts (multi-sorted algebras) in a language that has no support for them.

How hard would it be to extend C++ to really enable this style of programming? First, we need to introduce concepts as a new interface facility. For example, we can define:

```
concept SemiRegular : Assignable, DefaultConstructible {};
concept Regular : SemiRegular, EqualityComparable {};
concept InputIterator : Regular, Incrementable {
   SemiRegular value_type;
   Integral distance_type;
   const value_type& operator*();
};
value_type(InputIterator)
reduce(InputIterator first, InputIterator last, BinaryOperationWithIdentity op)
(value_type(InputIterator) == argument_type(BinaryOperationWithIdentity))
{
   if (first == last) return identity_element(op);
   value_type(InputIterator) result = *first;
   while (++first != last) result = op(result, *first);
   return result;
}
```

Generic functions are functions that take concepts as arguments and in addition to an argument list have a list of type constraints. Now full type checking can be done at the point

of definition without looking at the points of call, and full type-checking can be done at the points of call without looking at the body of the algorithm.

Sometimes we need multiple instances of the same concept. For example,

OutputIterator merge (InputIterator[1] first1, InputIterator[1] last1, InputIterator[2] first2, InputIterator[2] last2, OutputIterator result) (bool operator<(value_type(InputIterator[1]), value_type(InputIterator[2])), value_type(InputIterator[1]) == value_type(InputIterator[2]), output_type(OutputIterator) == value_type(InputIterator[2]));

Note that this merge is not as powerful as the STL merge. It cannot merge a list of *floats* and a vector of *doubles* into a deque of *ints*. STL algorithms will often do unexpected and, in my opinion, undesirable type conversions. If someone needs to merge *doubles* and *floats* into *ints*, he or she should use an explicit function object for asymmetric comparison and a special output iterator for conversion.

C++ provides two different abstraction mechanisms: object-orientedness and templates. Object-orientedness allows for exact interface definition and for run-time dispatch. But it cannot handle binary methods or multi-method dispatching, and its run-time binding is often inefficient. Templates handle richer interfaces and are resolved at compile-time. They can, however, cause a software engineering nightmare because of the lack of separation between interfaces and implementation. For example, I recently tried compiling a 10-line STL-based program using one of the most popular C++ compilers, and ran away in shock after getting several pages of incomprehensible error messages. And often one needs run-time dispatch that cannot be handled by templates. I do believe that introduction of concepts will unify both approaches and resolve both sets of limitations. And after all, it is possible to represent concepts as virtual tables that are extended by pointers to type descriptors: the virtual table for input iterator contains not just pointers to *operator** and *operator++*, but also pointers to the actual type of the iterator, its value type, and its distance type. And then one could introduce pointers to concepts and references to concepts!

Generic programming is a relatively young subdiscipline of computer science. I am happy to see that the small effort—started twenty years ago by Dave Musser, Deepak Kapur, Aaron Kershenbaum and me—led to a new generation of libraries such as BGL and MTL. And I have to congratulate Indiana University on acquiring one of the best generic programming teams in the world. I am sure they will do other amazing things!

Alexander Stepanov Palo Alto, California September, 2001¹

¹I would like to thank John Wilkinson, Mark Manasse, Marc Najork, and Jeremy Siek for many valuable suggestions.

Preface

The graph abstraction is a powerful problem-solving tool used to describe relationships between discrete objects. Many practical problems can be modeled in their essential form by graphs. Such problems appear in many domains: Internet packet routing, telephone network design, software build systems, Web search engines, molecular biology, automated road-trip planning, scientific computing, and so on. The power of the graph abstraction arises from the fact that the solution to a graph-theoretic problem can be used to solve problems in a wide variety of domains. For example, the problem of solving a maze and the problem of finding groups of Web pages that are mutually reachable can both be solved using depthfirst search, an important concept from graph theory. By concentrating on the essence of these problems—the graph model describing discrete objects and the relationships between them—graph theoreticians have created solutions to not just a handful of particular problems, but to entire families of problems.

Now a question arises. If graph theory is generally and broadly applicable to arbitrary problem domains, should not the software that implements graph algorithms be just as broadly applicable? Graph theory would seem to be an ideal area for software reuse. However, up until now the potential for reuse has been far from realized. Graph problems do not typically occur in a pure graph-theoretic form, but rather are embedded in larger domain-specific problems. As a result, the data to be modeled as a graph are often not explicitly represented as a graph but are instead encoded in some application-specific data structure. Even in the case where the application data are explicitly represented as a graph, the particular graph representation chosen by the programmer might not match the representation expected by a library that the programmer wants to use. Moreover, different applications may place different time and space requirements on the graph data structure.

This implies a serious problem for the graph library writer who wants to provide reusable software, for it is impossible to anticipate every possible data structure that might be needed and to write a different version of the graph algorithm specifically for each one. The current state of affairs is that graph algorithms are written in terms of whatever data structure is most convenient for the algorithm and users must convert their data structures to that format in order to use the algorithm. This is an inefficient undertaking, consuming programmer time and computational resources. Often, the cost is perceived not to be worthwhile, and the programmer instead chooses to rewrite the algorithm in terms of his or her own data structure. This approach is also time consuming and error prone, and will tend to lead to sub-optimal solutions since the application programmer may not be a graph algorithms expert.

Generic Programming

The Standard Template Library (STL) [40] was introduced in 1994 and was adopted shortly thereafter into the C++ Standard. The STL was a library of interchangeable components for solving many fundamental problems on sequences of elements. What set the STL apart from libraries that came before it was that each STL algorithm could work with a wide variety of sequential data structures: linked-lists, arrays, sets, and so on. The iterator abstraction provided an interface between containers and algorithms and the C++ template mechanism provided the needed flexibility to allow implementation without loss of efficiency. Each algorithm in the STL is a function template parameterized by the types of iterators upon which it operates. Any iterator that satisfies a minimal set of requirements can be used regardless of the data structure traversed by the iterator. The systematic approach used in the STL to construct abstractions and interchangeable components is called *generic programming*.

Generic programming lends itself well to solving the reusability problem for graph libraries. With generic programming, graph algorithms can be made much more flexible, allowing them to be easily used in a wide variety applications. Each graph algorithm is written not in terms of a specific data structure, but instead to a graph abstraction that can be easily implemented by many different data structures. Writing generic graph algorithms has the additional advantage of being more natural; the abstraction inherent in the pseudo-code description of an algorithm is retained in the generic function.

The Boost Graph Library (BGL) is the first C++ graph library to apply the notions of generic programming to the construction of graph algorithms.

Some BGL History

The Boost Graph Library began its life as the Generic Graph Component Library (GGCL), a software project at the Lab for Scientific Computing (LSC). The LSC, under the direction of Professor Andrew Lumsdaine, was an interdisciplinary laboratory dedicated to research in algorithms, software, tools, and run-time systems for high-performance computational science and engineering.² Special emphasis was put on developing industrial-strength, high-performance software using modern programming languages and techniques—most notably, generic programming.

Soon after the Standard Template Library was released, work began at the LSC to apply generic programming to scientific computing. The Matrix Template Library (MTL) was one of the first projects. Many of the lessons learned during construction of the MTL were applied to the design and implementation of the GGCL.

²The LSC has since evolved into the Open Systems Laboratory (OSL) *http://www.osl.iu.edu*. Although the name and location have changed, the research agenda remains the same.

An important class of linear algebra computations in scientific computing is that of sparse matrix computations, an area where graph algorithms play an important role. As the LSC was developing the sparse matrix capabilities of the MTL, the need for high-performance reusable (and generic) graph algorithms became apparent. However, none of the graph libraries available at the time (LEDA, GTL, Stanford GraphBase) were written using the generic programming style of the MTL and the STL, and hence did not fulfill the flexibility and high-performance requirements of the LSC. Other researchers were also expressing interest in a generic C++ graph library. During a meeting with Bjarne Stroustrup, we were introduced to several individuals at AT&T who needed such a library. Other early work in the area of generic graph algorithms included some codes written by Alexander Stepanov, as well as Dietmar Kühl's master's thesis.

With this in mind, and motivated by homework assignments in his algorithms class, Jeremy Siek began prototyping an interface and some graph classes in the spring of 1998. Lie-Quan Lee then developed the first version of the GGCL, which became his master's thesis project.

During the following year, the authors began collaborating with Alexander Stepanov and Matthew Austern. During this time, Stepanov's disjoint-sets-based connected components implementation was added to the GGCL, and work began on providing concept documentation for the GGCL, similar to Austern's STL documentation.

During this year the authors also became aware of Boost and were excited to find an organization interested in creating high-quality, open source C++ libraries. Boost included several people interested in generic graph algorithms, most notably Dietmar Kühl. Some discussions about generic interfaces for graph structures resulted in a revision of the GGCL that closely resembles the current Boost Graph Library interface.

On September 4, 2000, the GGCL passed the Boost formal review (managed by David Abrahams) and became the Boost Graph Library. The first release of the BGL was September 27, 2000. The BGL is not a "frozen" library. It continues to grow as new algorithms are contributed, and it continues to evolve to meet users' needs. We encourage readers to participate in the Boost group and help with extensions to the BGL.

What Is Boost?

Boost is an online community that encourages development and peer-review of free C++ libraries. The emphasis is on portable and high-quality libraries that work well with (and are in the same spirit as) the C++ Standard Library. Members of the community submit proposals (library designs and implementations) for review. The Boost community (led by a review manager) then reviews the library, provides feedback to the contributors, and finally renders a decision as to whether the library should be included in the Boost library collection. The libraries are available at the Boost Web site *http://www.boost.org*. In addition, the Boost mailing list provides an important forum for discussing library plans and for organizing collaboration.

Obtaining and Installing the BGL Software

The Boost Graph Library is available as part of the Boost library collection, which can be obtained in several different ways. The CD accompanying this book contains version 1.25.1 of the Boost library collection. In addition, releases of the Boost library collection can be obtained with your Web browser at *http://www.boost.org/boost_all.zip* for the Windows zip archive of the latest release and *http://www.boost.org/boost_all.tar.gz* for the UNIX archive of the latest release. The Boost libraries can also be downloaded via FTP at *ftp://boost.sourceforge.net/pub-/boost/release/*.

The zip archive of the Boost library collection can be unzipped by using WinZip or other similar tools. The UNIX "tar ball" can be expanded using the following command:

```
gunzip -cd boost_all.tar.gz | tar xvf -
```

Extracting the archive creates a directory whose name consists of the word **boost** and a version number. For example, extracting the Boost release 1.25.1 creates a directory **boost_1.25_1**. Under this top directory, are two principal subdirectories: **boost** and **libs**. The subdirectory **boost** contains the header files for all the libraries in the collection. The subdirectory **libs** contains a separate subdirectory for each library in the collection. These subdirectories contain library-specific source and documentation files. You can point your Web browser to **boost_1.25_1**/index.htm and navigate the whole Boost library collection.

All of the BGL header files are in the directory *boost/graph/*. However, other Boost header files are needed since BGL uses other Boost components. The HTML documentation is in *libs/graph/doc/* and the source code for the examples is in *libs/graph/example/*. Regression tests for BGL are in *libs/graph/test/*. The source files in *libs/graph/src/* implement the Graphviz file parsers and printers.

Except as described next, there are no compilation and build steps necessary to use BGL. All that is required is that the Boost header file directory be added to your compiler's include path. For example, using Windows 2000, if you have unzipped release 1.25.1 from *boost_all.zip* into the top level directory of your C drive, for Borland, GCC, and Metrowerks compilers add *-Ic:/boost_1_25_1* to the compiler command line, and for the Microsoft Visual C++ compiler add */I* "*c:/boost_1_25_1*". For IDEs, add *c:/boost_1_25_1* (or whatever you have renamed it to) to the include search paths using the appropriate dialog. Before using the BGL interface to LEDA or Stanford GraphBase, LEDA or GraphBase must be installed according to their installation instructions. To use the *read_graphviz()* functions (for reading AT&T Graphviz files), you must build and link to an additional library under *boost_1_25_1/libs/graph/src*.

The Boost Graph Library is written in ISO/IEC Standard C++ and compiles with most C++ compilers. For an up-to-date summary of the compatibility with a particular compiler, see the "Compiler Status" page at the Boost Web site *http://www.boost.org/status/compiler_status.html*.

How to Use This Book

This book is both a user guide and reference manual for the BGL. It is intended to allow the reader to begin using the BGL for real-life graph problems. This book should also be interesting for programmers who wish to learn more about generic programming. Although there are many books about how to use generic libraries (which in almost all cases means how to use the STL or Standard Library), there is very little available about how actually to build generic software. Yet generic programming is a vitally important new paradigm for software development. We hope that, by way of example, this book will show the reader how to do (and not simply use) generic programming and to apply and extend the generic programming paradigm beyond the basic container types and algorithms of the STL.

The third partner to the user guide and reference manual is the BGL code itself. The BGL code is not simply academic and instructional. It is intended to be used.

For students learning about graph algorithms and data structures, BGL provides a comprehensive graph algorithm framework. The student can concentrate on learning the important theory behind graph algorithms without becoming bogged down and distracted in too many implementation details.

For practicing programmers, BGL provides high-quality implementations of graph data structures and algorithms. Programmers will realize significant time saving from this reliability. Time that would have otherwise been spent developing (and debugging) complicated graph data structures and algorithms can now be spent in more productive pursuits. Moreover, the flexible interface to the BGL will allow programmers to apply graph algorithms in settings where a graph may only exist implicitly.

For the graph theoretician, this book makes a persuasive case for the use of generic programming for implementing graph-theoretic algorithms. Algorithms written using the BGL interface will have broad applicability and will be able to be reused in numerous settings.

We assume that the reader has a good grasp of C++. Since there are many sources where the reader can learn about C++, we do not try to teach it here (see the references at the end of the book—*The* C++ *Programming Language*, Special ed., by Bjarne Stroustrup [42] and C++ *Primer*, 3rd ed., by Josee Lajoie and Stanley B. Lippman [25] are our recommendations). We also assume some familiarity with the STL (see *STL Tutorial and Reference Guide* by David R. Musser, Gillmer J. Derge, and Atul Saini [34] and *Generic Programming and the STL* by Matthew Austern [3]). We do, however, present some of the more advanced C++ features used to implement generic libraries in general and the BGL in particular.

Some necessary graph theory concepts are introduced here, but not in great detail. For a detailed discussion of elementary graph theory see *Introduction to Algorithms* by T. H. Cormen, C. E. Leiserson, and R. L. Rivest [10].

Literate Programming

The program examples in this book are presented using the literate programming style developed by Donald Knuth. The literate programming style consists of writing source code and documentation together in the same file. A tool then automatically converts the file into both a pure source code file and into a documentation file with pretty-printed source code. The literate programming style makes it easier to ensure that the code examples in the book really compile and run and that they stay consistent with the text.

The source code for each example is broken up into *parts*. Parts can include references to other parts. For example, the following part labeled "Merge sort function definition" refers to the parts labeled "Divide the range in half and sort each half" and "Merge the two halves." An example often starts with a part that provides an outline for the entire computation, which is then followed by other parts that fill in the details. For example, the following function template is a generic implementation of the merge sort algorithm [10]. There are two steps in the algorithm, sorting each half of the range and then merging the two halves.

 \langle Merge sort function definition xxiia $\rangle \equiv$

Typically, the size of each part is limited to a few lines of code that carry out a specific task. The names for the parts are chosen to convey the essence of the task.

 \langle Divide the range in half and sort each half xxiib $\rangle \equiv$

```
RandomAccessIterator mid = first + (last - first)/2;
merge_sort(first, mid, cmp);
merge_sort(mid, last, cmp);
```

The *std::inplace_merge()* function does the main work of this algorithm, creating a single sorted range out of two sorted subranges.

 \langle Merge the two halves xxiic $\rangle \equiv$

std::inplace_merge(first, mid, last, cmp);

Parts are labeled with a descriptive name, along with the page number on which the part is defined. If more than one part is defined on a page, the definitions are distinguished by a letter.

Sometimes a file name is used for the label of a part. This means that the part is written out to a file. Many of the examples in the book are written out to files, and can be found in the *libs/graph/example/* directory of the Boost distribution. The following example shows the *merge_sort()* function being output to a header file. (merge-sort.hpp xxiii) ≡
#ifndef MERGE_SORT_HPP
#define MERGE_SORT_HPP

(Merge sort function definition xxiia)

#endif // MERGE_SORT_HPP

The Electronic Reference

An electronic version of the book is included on the accompanying CD, in the file *bgl-book.pdf*. The electronic version is searchable and is fully hyperlinked, making it a useful companion for the printed version. The hyperlinks include all internal references such as the literate programming "part" references as well as links to external Web pages.

Acknowledgments

We owe many debts of thanks to a number of individuals who both inspired and encouraged us in developing the BGL and in writing this book.

A most profound thanks goes to Alexander Stepanov and David Musser for their pioneering work in generic programming, for their continued encouragement of our work, and for contributions to the BGL. We especially thank David Musser for his careful proofreading of this book. Matthew Austern's work on documenting the concepts of the STL provided a foundation for creating the concepts in the BGL. We thank Dietmar Kühl for his work on generic graph algorithms and design patterns; especially for the property map abstraction. This work would not have been possible without the expressive power of Bjarne Stroustrup's C++ language.

Dave Abrahams, Jens Maurer, Dietmar Kühl, Beman Dawes, Gary Powell, Greg Colvin and the rest of the group at Boost provided valuable input to the BGL interface, numerous suggestions for improvement, and proofreads of this book. We also thank the following BGL users whose questions helped to motivate and improve BGL (as well as this book): Gordon Woodhull, Dave Longhorn, Joel Phillips, Edward Luke, and Stephen North.

Thanks to a number of individuals who reviewed the book during its development: Jan Christiaan van Winkel, David Musser, Beman Dawes, and Jeffrey Squyres.

A great thanks to our editor Deborah Lafferty; Kim Arney Mulcahy, Cherly Ferguson, and Marcy Barnes, the production coordinators; and the rest of the team at Addison–Wesley. It was a pleasure to work with them.

Our original work on the BGL was supported in part by NSF grant ACI-9982205. Parts of the BGL were completed while the third author was on sabbatical at Lawrence Berkeley National Laboratory (where the first two authors were occasional guests). All of the graph drawings in this book were produced using the dot program from the Graphviz package.

License

The BGL software is released under an open source "artistic" license. A copy of the BGL license is included with the source code in the LICENSE file.

The BGL may be used freely for both commercial and noncommercial use. The main restriction on BGL is that modified source code can only be redistributed if it is clearly marked as a nonstandard version of BGL. The preferred method for the distribution of BGL, and for submitting changes, is through the Boost Web site.

Chapter 3 A BGL Tutorial

As discussed in the previous chapter, *concepts* play a central role in generic programming. Concepts are the interface definitions that allow many different components to be used with the same algorithm. The Boost Graph Library defines a large collection of concepts that cover various aspects of working with a graph, such as traversing a graph or modifying its structure. In this chapter, we introduce these concepts and also provide some motivation for the choice of concepts in the BGL.

From the description of the generic programming process (see page 19), concepts are derived from the algorithms that are used to solve problems in particular domains. In this chapter we examine the problem of tracking file dependencies in a build system. For each subproblem, we examine generalizations that can be made to the solutions, with the goal of increasing the reusability (the genericity) of the solution. The result, at the end of the chapter, is a generic graph algorithm and its application to the file-dependency problem.

Along the way, we also cover some of the more mundane but necessary topics, such as how to create a graph object and fill in the vertices and edges.

3.1 File Dependencies

A common use of the graph abstraction is to represent dependencies. One common type of dependency that we programmers deal with on a routine basis is that of compilation dependencies between files in programs that we write. Information about these dependencies is used by programs such as *make*, or by IDEs such as Visual C++, to determine which files must be recompiled to generate a new version of a program (or, in general, of some target) after a change has been made to a source file.

Figure 3.1 shows a graph that has a vertex for each source file, object file, and library that is used in the *killerapp* program. An edge in the graph shows that a target depends on another target in some way (such as a dependency due to inclusion of a header file in a source file, or due to an object file being compiled from a source file).



Figure 3.1 A graph representing file dependencies.

Answers to many of the questions that arise in creating a build system such as make can be formulated in terms of the dependency graph. We might ask these questions:

- If all of the targets need to be made, in what order should that be accomplished?
- Are there any cycles in the dependencies? A dependency cycle is an error, and an appropriate message should be emitted.
- How many steps are required to make all of the targets? How many steps are required to make all of the targets if independent targets are made simultaneously in parallel (using a network of workstations or a multiprocessor, for example)?

In the following sections these questions are posed in graph terms, and graph algorithms are developed to provide solutions. The graph in Figure 3.1 is used in all of the examples.

3.2 Graph Setup

Before addressing these questions directly, we must first find a way to represent the filedependency graph of Figure 3.1 in memory. That is, we need to construct a BGL graph object.

Deciding Which Graph Class To Use

There are several BGL graph classes from which to choose. Since BGL algorithms are generic, they can also be used with any conforming user-defined graph class, but in this chapter we restrict our discussion to BGL graph classes. The principle BGL graph classes are the *adjacency_list* and *adjacency_matrix* classes. The *adjacency_list* class is a good choice for most situations, particularly for representing sparse graphs. The file-dependencies graph has only a few edges per vertex, so it is sparse. The *adjacency_matrix* class is a good choice for representing dense graphs, but a very bad choice for sparse graphs.

The *adjacency_list* class is used exclusively in this chapter. However, most of what is presented here also applies directly to the *adjacency_matrix* class because its interface is almost identical to that of the *adjacency_list* class. Here we use the same variant of *adjacency_list* as was used in $\S1.4.1$.

typedef adjacency_list<

listS, // Store out-edges of each vertex in a std::list vecS, // Store vertex set in a std::vector directedS // The file dependency graph is directed > file_dep_graph;

Constructing a Graph Using Edge Iterators

In §1.2.4 we showed how the *add_vertex()* and *add_edge()* functions can be used to create a graph. Those functions add vertices and edges one at a time, but in many cases one would like to add them all at once. To meet this need the *adjacency_list* graph class has a constructor that takes two iterators that define a range of edges. The edge iterators can be any InputIterator that dereference to a *std::pair* of integers (i, j) that represent an edge in the graph. The two integers i and j represent vertices where $0 \le i < |V|$ and $0 \le j < |V|$. The *n* and *m* parameters say how many vertices and edges will be in the graph. These parameters are optional, but providing them improves the speed of graph construction. The graph properties parameter *p* is attached to the graph object. The function prototype for the constructor that uses edge iterators is as follows:

template <typename EdgeIterator>
adjacency_list(EdgeIterator first, EdgeIterator last,
vertices_size_type n = 0, edges_size_type m = 0,
const GraphProperties& p = GraphProperties())

The following code demonstrates the use of the edge iterator constructor to create a graph. The *std::istream_iterator* is used to make an input iterator that reads the edges in from the file. The file contains the number of vertices in the graph, followed by pairs of numbers that specify the edges. The second default-constructed input iterator is a placeholder for the end of the input. The *std::istream_iterator* is passed directly into the constructor for the graph. std::ifstream file_in ("makefile-dependencies.dat ");
typedef graph_traits<file_dep_graph>::vertices_size_type size_type;
size_type n_vertices;
file_in >> n_vertices; // read in number of vertices
std::istream_iterator<std::pair<size_type, size_type> input_begin(file_in), input_end;
file_dep_graph g(input_begin, input_end, n_vertices);

Since the value type of the *std::istream_iterator* is *std::pair*, an input operator needs to be defined for *std::pair*.

```
namespace std {
  template <typename T>
  std::istream& operator>>(std::istream& in, std::pair<T,T>& p) {
    in >> p.first >> p.second;
    return in;
  }
}
```

3.3 Compilation Order

The first question that we address is that of specifying an order in which to build all of the targets. The primary consideration here is ensuring that before building a given target, all the targets that it depends on are already built. This is, in fact, the same problem as in §1.4.1, scheduling a set of errands.

3.3.1 Topological Sort via DFS

As mentioned in §1.4.2, a topological ordering can be computed using a depth-first search (DFS). To review, a DFS visits all of the vertices in a graph by starting at any vertex and then choosing an edge to follow. At the next vertex another edge is chosen to follow. This process continues until a dead end (a vertex with no out-edges that lead to a vertex not already discovered) is reached. The algorithm then backtracks to the last discovered vertex that is adjacent to a vertex that is not yet discovered. Once all vertices reachable from the starting vertex are explored, one of the remaining unexplored vertices is chosen and the search continues from there. The edges traversed during each of these separate searches form a *depth-first tree*; and all the searches form a *depth-first forest*. A depth-first forest for a given graph is not unique; there are typically several valid DFS forests for a graph because the order in which the adjacent vertices are visited is not specified. Each unique ordering creates a different DFS tree.

Two useful metrics in a DFS are the *discover time* and *finish time* of a vertex. Imagine that there is an integer counter that starts at zero. Every time a vertex is first visited, the value of the counter is recorded as the discover time for that vertex and the value of the counter is incremented. Likewise, once all of the vertices reachable from a given vertex have been

visited, then that vertex is finished. The current value of the counter is recorded as the finish time for that vertex and the counter is incremented. The discover time of a parent in a DFS tree is always earlier than the discover time of a child. Similarly, the finish time of a parent is always later than the finish time of a child. Figure 3.2 shows a depth-first search of the file dependency graph, with the tree edges marked with black lines and with the vertices labeled with their discover and finish times (written as discover/finish).



Figure 3.2 A depth-first search of the file dependency graph. Edges in the DFS tree are black and non-tree edges are gray. Each vertex is labeled with its discover and finish time.

The relationship between topological ordering and DFS can be explained by considering three different cases at the point in the DFS when an edge (u, v) is examined. For each case, the finish time of v is always earlier than the finish time of u. Thus, the finish time is simply the topological ordering (in reverse).

- 1. Vertex v is not yet discovered. This means that v will become a descendant of u and will therefore end up with a finish time earlier than u because DFS finishes all descendants of u before finishing u.
- 2. Vertex v was discovered in an earlier DFS tree. Therefore, the finish time of v must be earlier than that of u.
- 3. Vertex v was discovered earlier in the current DFS-tree. If this case occurs, the graph contains a cycle and a topological ordering of the graph is not possible. A *cycle* is a path of edges such that the first vertex and last vertex of the path are the same vertex.

The main part of the depth-first search is a recursive algorithm that calls itself on each adjacent vertex. We will create a function named *topo_sort_dfs()* that will implement a depth-first search modified to compute a topological ordering. This first version of the function will be a straightforward, nongeneric function. In the following sections we will make modifications that will finally result in a generic algorithm.

The parameters to *topo_sort_dfs()* include the graph, the starting vertex, a pointer to an array to record the topological order, and an array for recording which vertices have been visited. The *topo_order* pointer starts at the end of the array and then decrements to obtain the topological ordering from the reverse topological ordering. Note that *topo_order* is passed by reference so that the decrement made to it in each recursive call modifies the original object (if *topo_order* were instead passed by value, the decrement would happen instead to a copy of the original object).

The vertex_t type and edge_t types are the vertex and edge descriptors for the file_dep_graph.

typedef graph_traits<file_dep_graph>::vertex_descriptor vertex_t; typedef graph_traits<file_dep_graph>::edge_descriptor edge_t;

3.3.2 Marking Vertices Using External Properties

Each vertex should be visited only once during the search. To record whether a vertex has been visited, we can mark it by creating an array that stores the mark for each vertex. In general, we use the term *external property storage* to refer to the technique of storing vertex or edge properties (marks are one such property) in a data structure like an array or hash table that is separate from the graph object (i.e., that is *external* to the graph). Property values are looked up based on some key that can be easily obtained from a vertex or edge descriptor. In this example, we use a version of *adjacency_list* where the the vertex descriptors are integers from zero to *num_vertices(g) - 1*. As a result, the vertex descriptors themselves can be used as indexes into the mark array.

3.3.3 Accessing Adjacent Vertices

In the *topo_sort_dfs()* function we need to access all the vertices adjacent to the vertex *u*. The BGL concept AdjacencyGraph defines the interface for accessing adjacent vertices. The function *adjacent_vertices()* takes a vertex and graph object as arguments and returns a pair

of iterators whose value type is a vertex descriptor. The first iterator points to the first adjacent vertex, and the second iterator points past the end of the last adjacent vertex. The adjacent vertices are not necessarily ordered in any way. The type of the iterators is the *adjacency_iterator* type obtained from the *graph_traits* class. The reference section for *adjacency_list* (§14.1.1) reveals that the graph type we are using, *adjacency_list*, models the AdjacencyGraph concept. We may therefore correctly use the function *adjacent_vertices*() with our file dependency graph. The code for traversing the adjacent vertices in *topo_sort_dfs*() follows.

```
\langle For each adjacent vertex, make recursive call 47 \rangle \equiv
```

```
graph_traits<file_dep_graph>::adjacency_iterator vi, vi_end;
for (tie(vi, vi_end) = adjacent_vertices(u, g); vi != vi_end; ++vi)
if (mark[*vi] == 0)
topo_sort_dfs(g, *vi, topo_order, mark);
```

3.3.4 Traversing All the Vertices

One way to ensure that an ordering is obtained for every vertex in the graph (and not just those vertices reachable from a particular starting vertex) is to surround the call to *topo_sort_dfs()* with a loop through every vertex in the graph. The interface for traversing all the vertices in a graph is defined in the VertexListGraph concept. The *vertices()* function takes a graph object and returns a pair of vertex iterators. The loop through all the vertices and the creation of the mark array is encapsulated in a function called *topo_sort()*.

```
void topo_sort(const file_dep_graph& g, vertex_t* topo_order)
{
    std::vector<int> mark(num_vertices(g), 0);
    graph_traits<file_dep_graph>::vertex_iterator vi, vi_end;
    for (tie(vi, vi_end) = vertices(g); vi != vi_end; ++vi)
        if (mark[*vi] == 0)
            topo_sort_dfs(g, *vi, topo_order, &mark[0]);
}
```

To make the output from *topo_sort()* more user friendly, we need to convert the vertex integers to their associated target names. We have the list of target names stored in a file (in the order that matches the vertex number) so we read in this file and store the names in an array, which we then use when printing the names of the vertices.

```
std::vector<std::string> name(num_vertices(g));
std::ifstream name_in("makefile-target-names.dat");
graph_traits<file_dep_graph>::vertex_iterator vi, vi_end;
for (tie(vi, vi_end) = vertices(g); vi != vi_end; ++vi)
name_in >> name[*vi];
```

Now we create the order array to store the results and then apply the topological sort function.

```
std::vector<vertex_t> order(num_vertices(g));
topo_sort(g, &order[0] + num_vertices(g));
for (int i = 0; i < num_vertices(g); ++i)
std::cout << name[order[i]] << std::endl;</pre>
```

The output is

zag.cpp zig.cpp foo.cpp bar.cpp zow.h boz.h zig.o yow.h dax.h zag.o foo.o bar.o libfoobar.a libzigzag.a killerapp

3.4 Cyclic Dependencies

One important assumption in the last section is that the file dependency graph does not have any cycles. As stated in §3.3.1, a graph with cycles does not have a topological ordering. A well-formed makefile will have no cycles, but errors do occur, and our build system should be able to catch and report such errors.

Depth-first search can also be used for the problem of detecting cycles. If DFS is applied to a graph that has a cycle, then one of the branches of a DFS tree will loop back on itself. That is, there will be an edge from a vertex to one of its ancestors in the tree. This kind of edge is called a *back edge*. This occurrence can be detected if we change how we mark vertices. Instead of marking each vertex as visited or not visited, we use a three-way coloring scheme: white means undiscovered, gray means discovered but still searching descendants, and black means the vertex and all of its descendants have been discovered. Three-way coloring is useful for several graph algorithms, so the header file *boost/graph/properties.hpp* defines the following enumerated type.

enum default_color_type { white_color, gray_color, black_color };

A cycle in the graph is identified by an adjacent vertex that is gray, meaning that an edge loops back to an ancestor. The following code is a version of DFS instrumented to detect cycles.

bool has_cycle_dfs(const file_dep_graph& g, vertex_t u, default_color_type* color)
{
 color[u] = gray_color;
 graph_traits<file_dep_graph>::adjacency_iterator vi, vi_end;
 for (tie(vi, vi_end) = adjacent_vertices(u, g); vi != vi_end; ++vi)
 if (color[*vi] == white_color)
 if (has_cycle_dfs(g, *vi, color))
 return true; // cycle detected, return immediately
 else if (color[*vi] == gray_color) // *vi is an ancestor!
 return true;
 color[u] = black_color;
 return false;
}

As with the topological sort, in the *has_cycle()* function the recursive DFS function call is placed inside of a loop through all of the vertices so that we catch all of the DFS trees in the graph.

```
bool has_cycle(const file_dep_graph& g)
{
   std::vector<default_color_type> color(num_vertices(g), white_color);
   graph_traits<file_dep_graph>::vertex_iterator vi, vi_end;
   for (tie(vi, vi_end) = vertices(g); vi != vi_end; ++vi)
      if (color[*vi] == white_color)
      if (has_cycle_dfs(g, *vi, &color[0]))
        return true;
   return false;
}
```

3.5 Toward a Generic DFS: Visitors

At this point we have completed two functions, *topo_sort()* and *has_cycle()*, each of which is implemented using depth-first search, although in slightly different ways. However, the fundamental similarities between the two functions provide an excellent opportunity for code reuse. It would be much better if we had a single generic algorithm for depth-first search that expresses the commonality between *topo_sort()* and *has_cycle()* and then used parameters to customize the DFS for each of the different problems.

The design of the STL gives us a hint for how to create a suitably parameterized DFS algorithm. Many of the STL algorithms can be customized by providing a user-defined function object. In the same way, we would like to parameterize DFS in such a way that *topo_sort()* and *has_cycle()* can be realized by passing in a function object.

Unfortunately, the situation here is a little more complicated than in typical STL algorithms. In particular, there are several different locations in the DFS algorithm where customized actions must occur. For instance, the *topo_sort()* function records the ordering at the bottom of the recursive function, whereas the *has_cycle()* function needs to insert an operation inside the loop that examines the adjacent vertices.

The solution to this problem is to use a function object with more than one callback member function. Instead of a single *operator()* function, we use a class with several member functions that are called at different locations (we refer to these places as *event points*). This kind of function object is called an *algorithm visitor*. The DFS visitor will have five member functions: *discover_vertex()*, *tree_edge()*, *back_edge()*, *forward_or_cross_edge()*, and *finish_vertex()*. Also, instead of iterating over the adjacent vertices, we iterator over out-edges to allow passing edge descriptors to the visitor functions and thereby provide more information to the user-defined visitor. This code for a DFS function has a template parameter for a visitor:

```
template <typename Visitor>
void dfs_v1 (const file_dep_graph& g, vertex_t u, default_color_type* color, Visitor vis)
ł
  color[u] = gray_color;
  vis.discover_vertex(u, g);
  graph_traits<file_dep_graph>::out_edge_iterator ei, ei_end;
 for (tie(ei, ei_end) = out_edges(u, g); ei != ei_end; ++ei)
    if (color[target(*ei, g)] == white_color) {
      vis.tree_edge(*ei, g);
      dfs_v1(g, target(*ei, g), color, vis);
    } else if (color[target(*ei, g)] == gray_color)
      vis.back_edge(*ei, g);
    else
      vis.forward_or_cross_edge(*ei, g);
  }
  color[u] = black_color;
  vis.finish_vertex(u, g);
}
template <typename Visitor>
void generic_dfs_v1 (const file_dep_graph& g, Visitor vis)
  std::vector<default_color_type> color(num_vertices(g), white_color);
  graph_traits<file_dep_graph>::vertex_iterator vi, vi_end;
 for (tie(vi, vi\_end) = vertices(g); vi != vi\_end; ++vi)
    if (color[*vi] == white\_color)
      dfs_v1(g, *vi, \&color[0], vis);
  }
}
```

The five member functions of the visitor provide the flexibility we need, but a user that only wants to add one action should not have to write four empty member functions. This is easily solved by creating a default visitor from which user-defined visitors can be derived.

```
struct default_dfs_visitor {
   template <typename V, typename G>
   void discover_vertex(V, const G&) {
   template <typename E, typename G>
   void tree_edge(E, const G&) {
    template <typename E, typename G>
   void back_edge(E, const G&) {
    template <typename E, typename G>
   void forward_or_cross_edge(E, const G&) {
    template <typename V, typename G>
   void finish_vertex(V, const G&) {
    };
```

To demonstrate that this generic DFS can solve our problems, we reimplement the *topo_sort()* and *has_cycle()* functions. First we need to create a visitor that records the topological ordering on the "finish vertex" event point. The code for this visitor follows.

```
struct topo_visitor : public default_dfs_visitor {
  topo_visitor (vertex_t*& order) : topo_order(order) { }
  void finish_vertex(vertex_t u, const file_dep_graph&) {
    *--topo_order = u;
  }
  vertex_t*& topo_order;
};
```

Only two lines of code are required in the body of *topo_sort()* when implemented using generic DFS. One line creates the visitor object and one line calls the generic DFS.

```
void topo_sort(const file_dep_graph& g, vertex_t* topo_order)
{
   topo_visitor vis(topo_order);
   generic_dfs_v1(g, vis);
}
```

To reimplement the *has_cycle()* function, we use a visitor that records that the graph has a cycle whenever the back edge event point occurs.

```
struct cycle_detector : public default_dfs_visitor {
   cycle_detector(bool& cycle) : has_cycle(cycle) { }
   void back_edge(edge_t, const file_dep_graph&) {
      has_cycle = true;
   }
   bool& has_cycle;
};
```

The new *has_cycle()* function creates a cycle detector object and passes it to the generic DFS.

```
bool has_cycle(const file_dep_graph& g)
{
    bool has_cycle = false;
    cycle_detector vis(has_cycle);
    generic_dfs_v1(g, vis);
    return has_cycle;
}
```

3.6 Graph Setup: Internal Properties

Before addressing the next question about file dependencies, we are going to take some time out to switch to a different graph type. In the previous sections we used arrays to store information such as vertex names. When vertex or edge properties have the same lifetime as the graph object, it can be more convenient to have the properties somehow embedded in the graph itself (we call these *internal properties*). If you were writing your own graph class you might add data members for these properties to a vertex or edge struct.

The *adjacency_list* class has template parameters that allow arbitrary properties to be attached to the vertices and edge: the *VertexProperties* and *EdgeProperties* parameters. These template parameters expect the argument types to be the *property*<*Tag*, *T*> class, where *Tag* is a type that specifies the property and *T* gives the type of the property object. There are a number of predefined property tags (see §15.2.3) such as *vertex_name_t* and *edge_weight_t*. For example, to attach a *std::string* to each vertex use the following property type:

```
property<vertex_name_t, std::string>
```

If the predefined property tags do not meet your needs, you can create a new one. One way to do this is to define an enumerated type named *vertex_xxx_t* or *edge_xxx_t* that contains an enum value with the same name minus the *t* and give the enum value a unique number. Then use *BOOST_INSTALL_PROPERTY* to create the required specializations of the *property_kind* and *property_num* traits classes.¹ Here we create compile-time cost property that we will use in the next section to compute the total compile time.

```
namespace boost {
  enum vertex_compile_cost_t { vertex_compile_cost = 111 }; // a unique #
  BOOST_INSTALL_PROPERTY (vertex, compile_cost);
}
```

The *property* class has an optional third parameter that can be used to nest multiple *property* classes thereby attaching multiple properties to each vertex or edge. Here we create a new typedef for the graph, this time adding two vertex properties and an edge property.

¹Defining new property tags would be much simpler if more C++ compilers were standards conformant.

```
typedef adjacency_list<
listS, // Store out-edges of each vertex in a std::list
listS, // Store vertex set in a std::list
directedS, // The file dependency graph is directed
// vertex properties
property<vertex_name_t, std::string,
property<vertex_compile_cost_t, float,
property<vertex_distance_t, float,
property<vertex_color_t, default_color_type> >>>,
// an edge property
property<edge_weight_t, float>
> file_dep_graph2;
```

We have also changed the second template argument to *adjacency_list* from *vecS* to *listS*. This has some important implications. If we were to remove a vertex from the graph it would happen in constant time (with *vecS* the vertex removal time is linear in the number of vertices and edges). On the down side, the vertex descriptor type is no longer an integer, so storing properties in arrays and using the vertex as an offset will no longer work. However, the separate storage is no longer needed because we now have the vertex properties stored in the graph.

In §1.2.2 we introduced the notion of a property map. To review, a property map is an object that can be used to map from a key (such as a vertex) to a value (such as a vertex name). When properties have been specified for an *adjacency_list* (as we have just done), property maps for these properties can be obtained using the PropertyGraph interface. The following code shows an example of obtaining two property maps: one for vertex names and another for compile-time cost. The *property_map* traits class provides the type of the property map.

typedef property_map<file_dep_graph2, vertex_name_t>::type name_map_t; typedef property_map<file_dep_graph2, vertex_compile_cost_t>::type compile_cost_map_t; typedef property_map<file_dep_graph2, vertex_distance_t>::type distance_map_t; typedef property_map<file_dep_graph2, vertex_color_t>::type color_map_t;

The get() function returns a property map object.

name_map_t name_map = get(vertex_name, g); compile_cost_map_t compile_cost_map = get(vertex_compile_cost, g); distance_map_t distance_map = get(vertex_distance, g); color_map_t color_map = get(vertex_color, g);

There will be another file containing the estimated compile time for each makefile target. We read this file using a *std::ifstream* and write the properties into the graph using the property maps, *name_map* and *compile_cost_map*. These property maps are models of *LvaluePropertyMap* so they have an *operator[]()* that maps from vertex descriptors to a reference to the appriopriate vertex property object.

```
std::ifstream name_in( "makefile-target-names.dat ");
std::ifstream compile_cost_in( "target-compile-costs.dat ");
graph_traits<file_dep_graph2>::vertex_iterator vi, vi_end;
for (tie(vi, vi_end) = vertices(g); vi != vi_end; ++vi) {
    name_in >> name_map[*vi];
    compile_cost_in >> compile_cost_map[*vi];
}
```

In the following sections we will modify the topological sort and DFS functions to use the property map interface to access vertex properties instead of hard-coding access with a pointer to an array.

3.7 Compilation Time

The next questions we need to answer are, "How long will a compile take?" and "How long will a compile take on a parallel computer?" The first question is easy to answer. We simply sum the compile time for all the vertices in the graph. Just for fun, we do this computation using the *std::accumulate* function. To use this function we need iterators that, when dereferenced, yield the compile cost for the vertex. The vertex iterators of the graph do not provide this capability. When dereferenced, they yield vertex descriptors. Instead, we use the *graph_property_iter_range* class (see §16.8) to generate the appropriate iterators.

```
graph_property_iter_range<file_dep_graph2, vertex_compile_cost_t>::iterator ci, ci_end;

tie(ci, ci_end) = get_property_iter_range(g, vertex_compile_cost);

std::cout << "total (sequential) compile time: "

<< std::accumulate(ci, ci_end, 0.0) << std::endl;
```

The output of the code sequence is

total (sequential) compile time: 21.3

Now suppose we have a parallel super computer with hundreds of processors. If there are build targets that do not depend on each other, then they can be compiled at the same time on different processors. How long will the compile take now? To answer this, we need to determine the critical path through the file dependency graph. Or, to put it another way, we need to find the longest path through the graph.

The black lines in Figure 3.3 show the file dependency of *libfoobar.a*. Suppose that we have already determined when *bar.o* and *foo.o* will finish compiling. Then the compile time for *libfoobar.a* will be the longer of the times for *bar.o* and *foo.o* plus the cost for linking them together to form the library file.

Now that we know how to compute the "distance" for each vertex, in what order should we go through the vertices? Certainly if there is an edge (u, v) in the graph, then we better compute the distance for u before v because computing the distance to v requires the distance to u. This should sound familiar. We need to consider the vertices in topological order.



Figure 3.3 Compile time contributions to *libfoobar.a*.

3.8 A Generic Topological Sort and DFS

Due to the change in graph type (from *file_dep_graph* to *file_dep_graph2*) we can no longer use the *topo_sort()* function that we developed in §3.4. Not only does the graph type not match, but also the *color* array used inside of *generic_dfs_v1()* relies on the fact that vertex descriptors are integers (which is not true for *file_dep_graph2*). These problems give us an opportunity to create an even more generic version of topological sort and the underlying DFS. We parameterize the *topo_sort()* function in the following way.

- The specific type *file_dep_graph* is replaced by the template parameter *Graph*. Merely changing to a template parameter does not help us unless there is a standard interface shared by all the graph types that we wish to use with the algorithm. This is where the BGL graph traversal concepts come in. For *topo_sort()* we need a graph type that models the VertexListGraph and IncidenceGraph concepts.
- Using a *vertex_t** for the ordering output is overly restrictive. A more generalized way to output a sequence of elements is to use an output iterator, just as the algorithms in the C++ Standard Library do. This gives the user much more options in terms of where to store the results.
- We need to add a parameter for the color map. To make this as general as possible, we only want to require what is *essential*. In this case, the *topo_sort()* function needs to be able to map from a vertex descriptor to a marker object for that vertex. The Boost Property Map Library (see Chapter 15) defines a minimalistic interface for performing

this mapping. Here we use the LvaluePropertyMap interface. The internal *color_map* that we obtained from the graph in §3.6 implements the LvaluePropertyMap interface, as does the color array we used in §3.3.4. A pointer to an array of color markers can be used as a property map because there are function overloads in *boost/property_map.hpp* that adapt pointers to satisfy the LvaluePropertyMap interface.

The following is the implementation of our generic $topo_sort()$. The $topo_visitor$ and $generic_dfs_v2()$ are discussed next.

```
template <typename Graph, typename OutputIterator, typename ColorMap>
void topo_sort(const Graph& g, OutputIterator topo_order, ColorMap color)
{
    topo_visitor<OutputIterator> vis(topo_order);
    generic_dfs_v2(g, vis, color);
}
```

The *topo_visitor* class is now a class template to accommodate the output iterator. Instead of decrementing, we now increment the output iterator (decrementing an output iterator is not allowed). To get the same reversal behavior as in the first version of *topo_sort()*, the user can pass in a reverse iterator or something like a front insert iterator for a list.

```
template <typename OutputIterator>
struct topo_visitor : public default_dfs_visitor {
   topo_visitor (OutputIterator& order) : topo_order(order) { }
   template <typename Graph>
   void finish_vertex(typename graph_traits<Graph>::vertex_descriptor u, const Graph&)
   { *topo_order++ = u; }
   OutputIterator& topo_order;
};
```

The generic DFS changes in a similar fashion, with the graph type and color map becoming parameterized. In addition, we do not *a priori* know the color type, so we must get the color type by asking the *ColorMap* for its value type (though the *property_traits* class). Instead of using constants such as *white_color*, we use the color functions defined in *color_traits*.

```
template <typename Graph, typename Visitor, typename ColorMap>
void generic_dfs_v2(const Graph& g, Visitor vis, ColorMap color)
{
    typedef color_traits<typename property_traits<ColorMap>::value_type> ColorT;
    typename graph_traits<Graph>::vertex_iterator vi, vi_end;
    for (tie(vi, vi_end) = vertices(g); vi != vi_end; ++vi)
        color[*vi] = ColorT::white();
    for (tie(vi, vi_end) = vertices(g); vi != vi_end; ++vi)
        if (color[*vi] == ColorT::white())
        dfs_v2(g, *vi, color, vis);
}
```

The logic from the dfs_vI does not need to change; however, there are a few small changes required due to making the graph type parameterized. Instead of hard-coding *vertex_t* as the vertex descriptor type, we extract the appropriate vertex descriptor from the graph type using *graph_traits*. The fully generic DFS function follows. This function is essentially the same as the BGL *depth_first_visit()*.

```
template <typename Graph, typename ColorMap, typename Visitor>
void dfs_v2 (const Graph& g,
  typename graph_traits<Graph>::vertex_descriptor u,
  ColorMap color, Visitor vis)
{
  typedef typename property_traits<ColorMap>::value_type color_type;
  typedef color_traits<color_type> ColorT;
  color[u] = ColorT::gray();
  vis.discover_vertex(u, g);
  typename graph_traits<Graph>::out_edge_iterator ei, ei_end;
 for (tie(ei, ei_end) = out_edges(u, g); ei != ei_end; ++ei)
    if (color[target(*ei, g)] == ColorT::white()) {
      vis.tree_edge(*ei, g);
      dfs_v2(g, target(*ei, g), color, vis);
    } else if (color[target(*ei, g)] == ColorT::gray())
      vis.back_edge(*ei, g);
    else
      vis.forward_or_cross_edge(*ei, g);
  color[u] = ColorT::black();
  vis.finish_vertex(u, g);
}
```

The real BGL *depth_first_search*() and *topological_sort*() functions are quite similar to the generic functions that we developed in this section. We give a detailed example of using the BGL *depth_first_search*() function in $\S4.2$, and the documentation for *depth_first_search*() is in $\S13.2.3$. The documentation for *topological_sort*() is in $\S13.2.5$.

3.9 Parallel Compilation Time

Now that we have a generic topological sort and DFS, we are ready to solve the problem of finding how long the compilation will take on a parallel computer. First, we perform a topological sort, storing the results in the *topo_order* vector. We pass the reverse iterator of the vector into *topo_sort()* so that we end up with the topological order (and not the reverse topological order).

```
std::vector<vertex_t> topo_order(num_vertices(g));
topo_sort(g, topo_order.rbegin(), color_map);
```

Before calculating the compile times we need to set up the distance map (which we are using to store the compile time totals). For vertices that have no incoming edges (we call these source vertices), we initialize their distance to zero because compilation of these makefile targets can start right away. All other vertices are given a distance of infinity. We find the source vertices by marking all vertices that have incoming edges.

```
graph_traits<file_dep_graph2>::vertex_iterator i, i_end;
graph_traits<file_dep_graph2>::adjacency_iterator vi, vi_end;
```

```
// find source vertices with zero in-degree by marking all vertices with incoming edges
for (tie(i, i_end) = vertices(g); i != i_end; ++i)
    color_map[*i] = white_color;
for (tie(i, i_end) = vertices(g); i != i_end; ++i)
    for (tie(vi, vi_end) = adjacent_vertices(*i, g); vi != vi_end; ++vi)
        color_map[*vi] = black_color;
// initialize distances to zero, or for source vertices to the compile cost
for (tie(i, i_end) = vertices(g); i != i_end; ++i)
    if (color_map[*i] == white_color)
    distance_map[*i] = compile_cost_map[*i];
else
    distance_map[*i] = 0;
```

Now we are ready to compute the distances. We go through all of the vertices stored in *topo_order*, and for each one we update the distance (total compile time) for each adjacent vertex. What we are doing here is somewhat different than what was described earlier. Before, we talked about each vertex looking "up" the graph to compute its distance. Here, we have reformulated the computation so that instead we are pushing distances "down" the graph. The reason for this change is that looking "up" the graph requires access to in-edges, which our graph type does not provide.

```
std::vector<vertex_t>::iterator ui;
for (ui = topo_order.begin(); ui != topo_order.end(); ++ui) {
    vertex_t u = *ui;
    for (tie(vi, vi_end) = adjacent_vertices(u, g); vi != vi_end; ++vi)
        if (distance_map[*vi] < distance_map[u] + compile_cost_map[*vi])
            distance_map[*vi] = distance_map[u] + compile_cost_map[*vi];
}
```

The maximum distance value from among all the vertices tells us the total parallel compile time. Again we use *graph_property_iter_range* to create property iterators over vertex distances. The *std::max_element()* function does the work of locating the maximum.

The output is

total (parallel) compile time: 11.9

Figure 3.4 shows two numbers for each makefile target: the compile cost for the target and the time at which the target will finish compiling during a parallel compile.



Figure 3.4 For each vertex there are two numbers: compile cost and accumulated compile time. The critical path consists of black lines.

3.10 Summary

In this chapter we have applied BGL to answer several questions that would come up in constructing a software build system: In what order should targets be built? Are there any cyclic dependencies? How long will compilation take? In answering these questions we looked at topological ordering of a directed graph and how this can be computed via a depth-first search.

To implement the solutions we used the BGL *adjacency_list* to represent the file dependency graph. We wrote straightforward implementations of topological sort and cycle detection. We then identified common pieces of code and factored them out into a generic implementation of depth-first search. We used algorithm visitors to parameterize the DFS and then wrote specific visitors to implement the topological sort and the cycle detection.

We then looked at using a different variation of the *adjacency_list* class that allowed properties such as vertex name and compile cost to be attached to the vertices of the graph. We then further generalized the generic DFS by parameterizing the graph type and the property access method. The chapter finished with an application of the generic topological sort and DFS to compute the time it would take to compile all the targets on a parallel computer.

Index

, (comma), 40 . (period), 40 ; (semicolon), 73

A

abstract data types (ADTs), 19 accumulate function, 26-27 Adaptor(s) basic description of, 13-14 implementing, 123–126 pattern, 119 add_edge function, 9, 17, 43, 84, 121, 152-153, 226 EdgeMutablePropertyGraph concept and, 157 performance guidelines and, 128 undirected graphs and, 141 AdditiveAbelianGroup class, 20-21 add_vertex function, 9, 43, 120, 152, 157, 128, 225 AdjacencyGraph concept, 46–47, 114, 115, 146-149 adjacency_graph_tag function, 124 adjacency_iterator function, 47, 146 adjacency_list class, 11-12, 13 Bacon numbers and, 63, 65 basic description of, 43

boost namespace and, 37-39 compilation order and, 37, 46 implicit graphs and, 114 interfacing with other graph libraries and, 119 internal properties and, 52–53 maximum flow and, 107 minimum-spanning-tree problem and, 92 performance guidelines and, 127, 128, 130, 132 shortest-path problems and, 84 template parameters, 52 using topological sort with, 17–18 adjacency_list.hpp, 17, 216, 246 adjacency_matrix class, 11-12, 43 associated types, 238-239 basic description of, 235-242 member functions, 239-240 nonmember functions, 240-242 template parameters, 238 type requirements, 238 adjacency_matrix.hpp, 237 adjacency_vertices function, 46-47, 146 implicit graphs and, 114, 115 maximum flow and, 109 adjacent iterators, 7 ADTs (abstract data types), 19

advance_dispatch function, 33 advance function, 33 Algorithms. See also Algorithms (listed by name) basic description of, 13–18, 61–74 generic, 13-18 Koenig lookup and, 39 Algorithms (listed by name). See also Algorithms bellman_ford_shortest_paths algorithm, 40, 76-82, 162, 182-186 breadth_first_search algorithm, 11, 39, 61-67, 158-159, 165-169 depth first search algorithm, 13, 18, 44-46, 57, 67-75, 98, 160-161, 170-175 Dijkstra's shortest-path algorithm, 76, 81-88, 161, 179-181, 277 Edmunds-Karp algorithm, 105, 109 Ford-Fulkerson algorithm, 105 Kruskal's algorithm, 90-93, 95, 189-192 Prim's algorithm, 89, 90, 94–96 push-relabel algorithm, 105 ANSI C, 262 archetype class, 36 array_traits class, 31-33 array traits, for pointer types, 32-33 Array type, 30 Assignable concept, 37, 28–29, 143 Associated types, 28–34, 143, 205 adjacency_list class, 216-217 adjacency_matrix class, 238-239 edge_list class, 251-252 filtered raph class, 259–260 graph property iter range class, 298 iterator_property_map class, 284 LEDA Graph class, 268–269

property_map class, 249 reverse_graph class, 253–255 associative_property_map adaptor, 103 Austern, Matthew, 28

B

back_edge function, 50, 67, 160 back_edge_recorder class, 70 back_edges vector, 71 backward edge, 106 Bacon, Kevin, 62–67 bacon_number array, 66 bacon_number_recorder, 66 Bacon numbers basic description of, 61–67 graph setup and, 63-65 input files and, 63-65 bar.o. 54 Base classes, 20, 21 parameter, 37 basic block, 69 BCCL (Boost Concept Checking Library), 35, 36 bellman_ford.cpp, 185–186 bellman ford shortest paths algorithm, 40, 162 basic description of, 76-82, 182-186 named parameters, 184 parameters, 183 time complexity and, 185 BellmanFordVisitor concept, 161–162 BFS (breadth-first search), 11, 39. See also breadth_first_search algorithm Bacon numbers and, 65-67 basic description of, 61-67 visitor concepts and, 158-159 bfs_name_printer class, 11

INDEX

BFSVisitor interface, 66, 158–159 bgl named params class, 40 BidirectionalGraph concept, 69, 72, 124, 145 - 146bidirectional graph tag class, 124 Binary method problem, 23–24 boost::array, 78 Boost Concept Checking Library (BCCL), 35,36 boost::forward_iterator_helper, 114 BOOST INSTALL PROPERTY, 52 Boost namespace adjacency_list class and, 37-38 basic description of, 37-39 boost:: prefix, 37-38 Boost Property Map Library, 55-56, 79, 80 Boost Tokenizer Library, 63 breadth-first search (BFS), 11, 39. See also breadth_first_search algorithm Bacon numbers and, 65-67 basic description of, 61-67 visitor concepts and, 158-159 breadth_first_search algorithm, 11, 39. See also breadth-first search (BFS) Bacon numbers and, 65-67 basic description of, 61-67, 165-169 named parameters, 167 parameters, 166 preconditions, 167 visitor concepts and, 158-159 Breadth-first tree, 61

C

C++ (high-level language) associated types and, 28 binary method problem and, 23–24 code bloat and, 23 concept checking and, 34–37

expressions, 28 generic programming in, 19-59 GNU, 127, 128 Koenig lookup and, 38–39 Monoid concept and, 291 named parameters and, 39-40 object-oriented programming and, 22-25 Standard, 22, 28, 55, 125 valid expressions and, 28 capacity_map parameter, 206 cap object, 108 cc-internet.dot, 98 Chessboard, Knight's tour problem for, 113 - 118Class(es). See also Classes (listed by name) abstract, 20 archetype, 36 auxiliary, 242-251, 289-298 basic description of, 13-14, 213-275 base, 20, 21 comparisons, 127-132 concept-checking, 35-36 nesting, 52 selecting, 43-44 typedefs nested in, 30-31 Classes (listed by name). See also adjacency_list class AdditiveAbelianGroup class, 20-21 adjacency_matrix class, 11-12, 43, 234-242 adjacency_vertices class, 109, 114, 115 archetype class, 36 array_traits class, 31-33 back_edge_recorder class, 70 bgl named params class, 40 ColorPoint class, 23-24 ColorPoint2 class, 24 color traits class, 56, 290

Classes, *continued* edge list class, 78-79 equivalence class, 98 filtered_graph class, 256-262 GRAPH class, 120, 123–126 graph_property_iter_range class, 58, 297-298 graph_traits class, 5-7, 33-34, 47, 57, 125-126, 142-148, 242-245 iterator_adaptor class, 125 iterator_property_map class, 79, 283–285 iterator_traits class, 29, 33, 278 johns_int_array class, 32 knights_tour_graph class, 114–117 make iterator property map class, 79 parallel_edge_traits class, 234 Point class, 24 property class, 52-53, 63, 231 property_kind class, 52 property_map class, 53, 156, 248–249 property num class, 52 property_traits class, 33, 56, 231, 278, 282 put_get_helper class, 286 sgb_vertex_id_map class, 122 sgb vertex name map class, 286 std::iterator_traits class, 29 clear_vertex function, 134, 153, 226 clock function, 128 Code. See also Source code files "bloat," 23 size, 23 Color array, 55 connected components and, 100, 102-104 maps, 55-56, 69, 100, 167-181, 196, 199, 207 markers, 56

three-way, 48 types, accessing, 56 ColorMap, 56, 175, 176–177, 181 color_map parameter, 56, 69, 167, 170, 172, 176, 178, 180–181, 196, 199, 207 ColorPoint class, 23–24 ColorPoint2 class, 24 color_traits class, 56, 290 ColorValue concept, 175, 177, 290 Comma, 40 compare function, 183 Compilation. See also Compilers dispatch of virtual functions during, 22 order, 44–48 time, 22, 54–55, 57–59 compile_cost_map, 53 Compilers. See also Compilation Koenig lookup and, 38–39 partial specialization and, 32-33 Complexity, time adjacency_list class and, 225-227 basic description of, 165 bellman_ford_shortest_paths function and, 185 breadth_first_search function and, 167 connected components function and, 196 depth_first_search function and, 172 depth_first_visit function and, 176 dijkstra shortest paths function and, 181 disjoint sets class and, 295 incremental components and, 202, 203, 204 johnson_all_pairs_shortest_paths function and, 188 kruskal minimum spanning tree function and, 191 prim_minimum_spanning tree function and, 194

strong_components function and, 200 topological sort function and, 177 Complexity guarantees, 28, 145–148, 152 - 157Component array, 103 component_index function, 204-206 Components, connected basic description of, 97-104 static, 195-201 strongly, 97, 102-104 incremental, 201-205 Internet connectivity and, 98-101 Web page links and, 102–104 Component vector, 100 compute_loop_extent function, 71 Concept(s) archetypes, 36-37 auxiliary, 289-298 checking, 34-37 covering, 36-37 definitions, 27-28 generic programming, 27-29 graph modification, 150–157 graph traversal, 137–150 notation for, 137 property map, 278-281 refinement of, 28, 138 use of the term, 19 visitor, 158-160 Connected components basic description of, 97-104 static, 195-201 strongly, 97, 102-104 incremental, 201-205 Internet connectivity and, 98-101 Web page links and, 102–104 connected_components function, 98, 100, 195-197

connected_components.hpp, 196 Constructors, 43-44 container_gen.cpp, 232, 233 Containers basic description of, 25 hash_map container, 187 CopyConstructible, 36, 143 Cross edge, 68 Cut, capacity of, 106 cycle_edge function, 159 Cycle(s). See also Cyclic dependencies basic description of, 45 detector objects, 51 makefiles and, 48 visitors and, 51-52 Cyclic dependencies, 42, 48-49. See also Cycles

D

Data structures, traversal through, 25 default_bfs_visitor, 66, 70 DefaultConstructible, 143 degree_size_type function, 144, 146 Delay array, 79 Dependencies cyclic, 42, 48-49 file, 41-42, 54-55 Depth-first forest, 44, 67 Depth-first search (DFS). See also depth first search algorithm basic description of, 67-74 connected components and, 98 cyclic dependencies and, 48-49 generic, 49-52 topological sorts and, 55-57 upstream, 72 depth_first_search algorithm, 13, 18, 57

Depth-first search (DFS) basic description of, 67–75, 170–175 connected components and, 98 topological sort and, 44-46 named parameters, 172 parameters, 172 time complexity and, 172 topological sort and, 44-46 visitor concepts and, 160-161 Depth-first tree, 44, 67 depth first visit function, 57, 67, 70, 72 basic description of, 175-176 parameters, 175 time complexity and, 176 Descriptors, 5-6, 144, 148, 227-229 DFS (depth-first search). See also depth_first_search algorithm basic description of, 67-74 connected components and, 98 cyclic dependencies and, 48-49 generic, 49-52 topological sorts and, 55-57 upstream, 72 DFSVisitor interface, 160-161 dfs_vl, 57 difference_type, 29 DiffType parameter, 251 dijkstra.cpp, 227 Dijkstra's shortest-path algorithm, 76, 161, 277 basic description of, 81-88, 179-181 named parameters, 179 parameters, 179 time complexity and, 181 DIMACS file format, 207, 211 directed category function, 34, 124, 244 Directed parameter, 12, 214, 238, 246 directed_tag function, 124

Directed version, 138 disconnecting_set iterator, 110 discover_time parameter, 44–45, 199 discover_vertex function, 50 distance_combine parameter, 184 distance_compare parameter, 184 distance_map parameter, 40, 183, 184, 187, 194 distance_zero parameter, 187 ds.find_set function, 201–205 ds.union_set function, 201–205

E

Edge(s) adding, 128-129 backward, 106 connectivity, 106-112 cross, 68 descriptors, 5-6, 144, 148 forward, 68, 106 iterators, 7, 43-44, 148, 251 parallel, 4 relaxation, 77–78 removing, 130 residual capacity of, 105, 206 saturated, 105 edge_descriptor, 144, 148 edge function, 149, 226 edge iterator, 148 EdgeIterator parameter, 251 edge_length_t tag, 266 edge_list class, 78-79 EdgeListGraph interface, 78, 92, 147–148 edge_list.hpp, 251 EdgeList parameter, 12 edge list template, 78–79 EdgeList type, 65, 127, 132 EdgeMutableGraph concept, 124, 152, 154 EdgeMutablePropertyGraph concept, 157 edge_parallel_category function, 34, 124, 244 EdgeProperties parameter, 12, 52 EdgeProperty parameter, 238 edges function, 8, 147 edges_size_type, 148-149 edge type, 124 edge_weight_t tag, 52, 92 edge xxx t tag, 52 Edmunds-Karp algorithm, 105, 109 edmunds_karp_max_flow function, 206-209 empty function, 290 entry vertex, 70 enum, 78 equal function, 23-24 EqualityComparable, 28–29, 143 equivalence class, 98 Equivalence relations, 98 erase_dispatch function, 234 erase function, 234 Errors, concept checking and, 34–37 Erdös, Paul, 62 Erdös number, 62 Event points, 50 examine_edge function, 159

F

f function, 40 File dependencies basic description of, 41–42 compilation time and, 54–55 file_dep_graph function, 55 file_dep_graph2 function, 55 filtered_graph, 13, 257–262 find_loops function, 69 find_with_full_path_compression function, 295 find_with_path_halving function, 295 Finish time, 44–45 finish vertex function, 50 First_Adj_Edge, 126 first argument, 27 first variable, 8 Flow functions, 105. See also Maximum-flow algorithms Flow networks, 105. See also Maximum-flow algorithms foo.o, 54 Ford, L. R., 105 Ford-Fulkerson algorithm, 105 for each function, 29 Forward edge, 68, 106 ForwardIterator, 25-26 forward_iterator_tag function, 33 forward or cross edge function, 50 Fulkerson, D. R., 105 Function(s). See also Functions (listed by name) auxiliary, 289-298 objects, user-defined, 49 preconditions for, 165 prototypes, 43, 164 virtual, function templates and, comparison of, 22 Functions (listed by name). See also Functions(s) add_edge function, 9, 17, 43, 84, 121, 128, 141, 152–153, 157, 226 add_vertex function, 9, 43, 120, 152, 157, 128, 225 adjacency graph tag function, 124 adjacency_iterator function, 47, 146 adjacent vertices function, 46-47, 146

Functions, continued advance_dispatch function, 33 advance function, 33 allow_parallel_edge_tag function, 124 back edge function, 50, 67, 160 bfs name printer function, 11 bidirectional_graph_tag function, 124 clear vertex function, 134, 153, 226 clock function, 128 compare function, 183 component index function, 204-206 compute_loop_extent function, 71 connected_components function, 98, 100, 195-197 cycle edge function, 159 degree_size_type function, 144, 146 depth_first_visit function, 57, 67, 70, 72, 175 - 176directed_category function, 34, 124, 244 directed tag function, 124 discover vertex function, 50 ds.find_set function, 201-205 ds.union set function, 201-205 edge function, 149, 226 edge parallel category function, 34, 124, 244 edges function, 8, 147 edmunds_karp_max_flow function, 206-209 empty function, 290 equal function, 23-24 erase dispatch function, 234 erase function, 234 examine edge function, 159 f function, 40 file dep graph function, 55 file_dep_graph2 function, 55 find loops function, 69

find_with_full_path_compression function, 295 find_with_path_halving function, 295 finish vertex function, 50 for each function, 29 forward_iterator_tag function, 33 forward_or_cross_edge function, 50 function requires function, 35 gb_new_edge function, 123 generic_dfs_vl function, 55 get function, 6, 79-80, 156, 248, 266, 279, 286 has_cycle function, 49-52 identity property map function, 16 in edges function, 8 incremental components function, 195, 199 insert function, 64 johnson_all_pairs_shortest_paths function, 186-188 kruskal_minimum_spanning_tree function, 91-93, 189-192 link function, 294 make_back_edge_recorder function, 71 max element function, 58 num edges function, 149 num vertices function, 38-39, 46, 115, 147, 162 operator function, 50, 158 out degree function, 144, 145 out_edges function, 8, 72, 125, 126, 138, 140, 143–145 prim minimum spanning tree function, 94-96, 192-195 print equal function, 24 print equal2 function, 24 print_graph function, 272 print trans delay function, 6

print_vertex_name function, 6 push dispatch function, 234 push function, 234, 290 push_relabel_flow function, 206 push relabel max flow function, 209-212 put function, 279, 280, 286 read_graphviz function, 84, 91 remove_edge function, 152-153, 226, 228 remove edge if function, 227 remove_in_edge_if function, 155 remove_out_edge_if function, 154 remove vertex function, 152, 225, 226 safe_sort function, 35 size function, 290 sort function, 10, 34 source function, 124, 125, 138, 140, 145, 147 std::accumulate function, 54 std::advance function, 33 std::back_inserter function, 92 std::for each function, 29 strong components function, 102–103, 198-201 Succ Adj Edge function, 125 sum function, 20, 21, 30, 31, 32 target function, 124, 125, 138, 140, 145, 149 tie function, 8, 296 top function, 290 topological_sort function, 13, 14-18, 44-46, 57, 119, 120-123, 176-177 topo_sort_dfs function, 46-47 topo sort function, 47, 49, 51, 55-58 traversal category function, 34, 124, 244 tree_edge function, 50, 160 union set function, 294

valid_position function, 114–115 vertex function, 225 vertex_index_map function, 16, 187, 191, 194, 196, 199, 207 vertex_list_graph_tag function, 124 vertices function, 8, 47, 296 vertices_size_type function, 124 visitor function, 11, 66 who_owes_who function, 231 function_requires function, 35

G

"Gang of Four" (GoF) Patterns Book, 10-11 gb_new_edge function, 123 g_dot, 84 Generalized pointers, 25. See also Iterators generic_dfs_vl function, 55 Generic programming (GP) boost namespace and, 37-39 concepts, 27-29, 34-37 in C++, 19-59 Koenig lookup and, 38–39 models, 27–29 named parameters and, 39-40 object-oriented programming and, comparison of, 22-25 the STL and, 25-27 get function, 6, 79-80, 156, 248, 266, 279, 286 GNU C++, 127, 128. See also C++ (high-level language) Goldberg, A. V., 105 GP (generic programming). See Generic programming (GP) Graph(s) adaptors, 13-14, 119, 123-126 directed, 3-4 implicit, 113-118

Graph(s), *continued* internal properties for, 52-54, 229-230 libraries, interfacing with, 119–126 modification, 9-10, 150-157 search, backtracking, 116–116 setup, 42-44, 52-54, 63-65 terminology, 3-4 traversal, 7-8, 24, 124, 137-150, 244 undirected, 4, 138-142 graph_archetypes.hpp, 37 GRAPH class, 120, 123-126 Graph concept, 142–153 graph_concepts.hpp, 35 graph.cpp, 126 Graph parameter, 55 GraphProperties parameter, 12 graph_property_iter_range class, 58, 297–298 GraphProperty parameter, 238 graph_traits class, 5-7, 33-34, 47, 57, 125-126, 142-148 basic description of, 242-245 category tags, 244 template parameters, 244 graph traits.hpp, 244-245 Graph type, 69, 84, 94, 119–120 GraphvizDigraph, 82, 84, 103 GraphvizGraph type, 82, 91, 98, 100 graphviz.hpp, 84, 91 Graphviz.org, 83 Guarantees, complexity, 28, 145–148, 152-157 Guidelines, performance basic description of, 127-134 graph class comparisons and, 127-132

H

Hamlitonian path, 113–114 has_cycle function, 49–52

hash_map container, 187 Heuristics path compression, 190, 293 union by rank, 189–190, 293 Hop, use of the term, 76 .hpp files adjacency_list.hpp, 17, 216, 246 adjacency_matrix.hpp, 237 connected_components.hpp, 196 edge_list.hpp, 251 graph_archetypes.hpp, 37 graph_concepts.hpp, 35 graph_traits.hpp, 244-245 graphviz.hpp, 84, 91 johnson_all_pairs_shortest_path.hpp, 186 kruskal_minimum_spanning_tree.hpp, 190 leda_graph.hpp, 120-121, 126, 243, 266 prim minimum spanning tree.hpp, 193 properties.hpp, 48-49, 248, 250 property.hpp, 256 property_iter_range.hpp, 298 property_map.hpp, 56, 283 push_relabel_max_flow.hpp, 209 reverse graph.hpp, 253 stanford_graph.hpp, 13, 122, 243 strong_components.hpp, 198 vector_as_graph.hpp, 15, 17, 272

I

identity_property_map function, 16 Implicit graphs backtracking graph search and, 116–117 basic description of, 113–118 Warnsdorff's heuristic and, 117–118 in_degree, 145 In-edge(s) basic description of, 4–5 BidirectionalGraph concept and, 145-146 iterators, 7-8 undirected graphs and, 138, 140 in_edges function, 8 IncidenceGraph concept, 143-145, 154 incremental_components.cpp, 202 incremental_components function, 195, 199 InputIterator, 25–26, 28–29 insert function, 64 Interface, use of the term, 5 interior_property_map.cpp, 232 Internet. See also Routers; Routing connectivity, connected components and, 98-101 Movie Database, 63 Internet Protocol (IP), 76 Invariants, 28 IP (Internet Protocol), 76 iterator_adaptor class, 125 iterator_category type, 29, 33 iterator_property_map adaptor, 100 iterator_property_map class, 79, 283-285 Iterators adjacency_list and, 227-229 basic description of, 25 categories of, 25-27 constructing graphs using, 43-44 iterator traits class, 29, 33, 278

J

johns_int_array class, 32 johnson_all_pairs_shortest_path.hpp, 186 johnson_all_pairs_shortest_paths function basic description of, 186–187 named parameters, 187–188 parameters, 187

K

Karzanov, A. V., 105 Keyword parameters, 39-40 Killerapp programs, 41 knights_adjacency_iterator, 115 knights_tour_graph class, 114–117 Knight's tour problem, 113–118 Knuth, Donald, 119 Koenig lookup, 38–39 Kruskal, J. B., 89, 90, 94, 95 kruskal.cpp, 191 kruskal_minimum_spanning_tree function basic description of, 91–93, 189–192 named parameters, 190-191 parameters, 190 time complexity and, 191 kruskal_minimum_spanning_tree.hpp, 190 Kruskal's algorithm, 90–93, 95, 189–192

L

last argument, 27 last variable, 8 LEDA graphs, 119–121, 123–126 basic description of, 13 graph adaptors and, 13 adaptors for, 267-272 leda_g, 120 leda_graph.hpp, 120-121, 126, 243, 266 LessThanComparable interface, 34, 36, 37 lexical cast, 91 libfoobar.a, 54-55 lib_jack, 38 lib_jill, 39 link function, 294 Link-state advertisement, 81 Lists, vectors of, using topological sorts with, 14-17 listS argument, 17, 53, 132, 233

Loop(s) basic description of, 69 finding, in program-control-flow graphs, 69–73 head, 69 self-, 4 termination, 87 loop_set, 73 LvaluePropertyMap interface, 53, 56, 66, 79

Μ

make_back_edge_recorder function, 71 Makefiles, 48, 59 make_iterator_property_map class, 79 make_leda_node_property_map, 121 max_element function, 58 Max-Flow Min-Cut Theorem, 106 Maximum-flow algorithms basic description of, 105–112, 206-213 edge connectivity and, 106-112 miles_span.cpp, 262 Minimum diconnected set, 106 Minimum-spanning-tree problem basic description of, 89–96, 189–195 Kruskal's algorithm and, 91-93 Prim's algorithm and, 94-96 Model, use of the term, 21, 28 MultiPassInputIterator, 146 Musser, D. R., 25 MutableBidirectionalGraph concept, 154 - 155MutableEdgeListGraph concept, 155 MutableIncidenceGraph concept, 154 my array, 30-31 Myers, Nathan, 30

N

Named parameters. See also Parameters basic description of, 39-40, 164 bellman_ford_shortest_paths function, 184 breadth first search function, 167 breadth_first_visit function, 170 connected components function, 196 depth_first_search function, 172 dijkstra shortest paths function, 179-181 edmunds_karp_max_flow function, 206-207 johnson_all_pairs_shortest_paths function, 187-188 kruskal_minimum_spanning_tree function, 190-191 prim_minimum_spanning tree function, 194 push_relabel_max_flow function, 210-211 strong_components function, 199-200 topological_sort function, 176-177 name map, 53 Namespaces boost namespace, 37-39 Koenig lookup and, 38-39 Nesting classes, 52 NextProperty parameter, 229 node_array, 121 node type, 124 Notation, 137 num edges function, 149 num vertices function, 38-39, 46, 115, 147, 162

0

OOP (object-oriented programming) generic programming and, comparison of, 22–25 Graph concept and, 142 polymorphism and, 19-21 operator function, 50, 158 OSPF (Open Shortest Path First) protocol, 82 out_degree function, 144, 145 Out-edge(s) adaptors, 125 basic description of, 4-5 iterators, 7-8, 126, 144 traversal, 132 out edge adaptor, 125 out_edge_iterator, 126, 144 out_edges function, 8, 72, 125, 126 complexity guarantees and, 145 IncidenceGraph concept and, 143–144 undirected graphs and, 138, 140

P

Packets, basic description of, 76
Parallel compilation time, 57–59. See also Compilation
parallel_edge_traits class, 234
Parameters. See also Named parameters; Parameters (listed by name); Template parameters
adjacency_list class, 216
adjacency_list_traits class, 246
adjacency_matrix class, 238
adjacency_matrix_traits class, 247
basic description of, 39–40, 164
bellman_ford_shortest_paths function, 183
breadth_first_search function, 166

breadth_first_visit function, 170 connected components function, 196 depth_first_search function, 172 depth_first_visit function, 175 dijkstra shortest paths function, 179 disjoint sets class, 294 edge_list class, 251 edmunds karp max flow function, 206 filtered graph class, 259 graph_property_iter_range class, 298 graph_traits class, 244 iterator_property_map class, 284 johnson_all_pairs_shortest_paths function, 187 kruskal_minimum_spanning_tree function, 190 LEDA Graph class template, 268 mutable_queue adaptor, 292 prim minimum spanning tree function, 193 property class, 250 property_map class, 249 property_traits class, 283 push_relabel_max_flow function, 210 reverse graph class, 253 strong components function, 199 topological_sort function, 176 Parameters (listed by name). See also Parameters Base parameter, 37 capacity_map parameter, 206 color_map parameter, 56, 69, 167, 170, 172, 176, 178, 180–181, 196, 199, 207 DiffType parameter, 251 Directed parameter, 12, 214, 238, 246

Parameters, continued discover time parameter, 44-45, 199 distance_combine parameter, 184 distance_compare parameter, 184 distance_map parameter, 40, 183, 184, 187, 194 distance_zero parameter, 188 EdgeIterator parameter, 251 EdgeList parameter, 12 EdgeProperties parameter, 12, 52 EdgeProperty parameter, 238 Graph parameter, 55 GraphProperties parameter, 12 GraphProperty parameter, 238 NextProperty parameter, 229 predecessor_map parameter, 85, 183, 184, 190, 207 residual_capacity_map parameter, 206 reverse edge map parameter, 207 root vertex parameter, 194 topo_sort_visitor parameter, 18 ValueType parameter, 251 VertexProperties parameter, 12, 52 VertexProperty parameter, 238 visitor parameter, 162, 184 weight_map parameter, 79, 184, 187, 194 Parent(s) array, 95 basic description of, 61, 67 maps, 85 minimum-spanning-tree problem and, 94-95 shortest-path problems and, 80 Parsers, 82-84 Partial specialization, providing array traits for pointer types with, 32-33

Path(s). See also Shortest-path problems basic description of, 75, 97 Hamlitonian, 113-114 compression heuristics, 190, 293 path cost, 87 Performance guidelines basic description of, 127-134 graph class comparisons and, 127 - 132Period (.), 40 Point class, 24 Pointer types, 29, 32–33 Polymorphism basic description of, 19, 20, 21 parametric, 21, 22 subtype, 20, 22 pop function, 290 POSIX, 128 Pred_Adj_Edge function, 125 predecessor map parameter, 85, 183, 184, 190, 207 Predecessors, basic description of, 61 Prim, R., 89 prim.cpp, 195 prim_minimum_spanning_tree function basic description of, 94-96, 192-195 named parameters, 194 parameters, 193 time complexity and, 194 prim minimum spanning tree.hpp, 193 Prim's algorithm, 89, 90, 94-96 print equal function, 24 print_equal2 function, 24 print graph function, 272 print trans delay function, 6 print_vertex_name function, 6 Program-control-flow graphs, 69-73

Properties. See also Property maps; Property tags basic description of, 5 custom, 230 external storage of, 46 internal, 52-54, 229-230 marking vertices using, 46 properties.hpp, 48-49, 248, 250 property class, 52-53, 63, 231 PropertyGraph interface, 53, 155–156 property.hpp, 256 property_iter_range.hpp, 298 property_kind class, 52 Property map(s), 53, 103basic description of, 6–7 classes, 281-285 concepts, 278-281 creating your own, 283-287 implemented with std::map, 287 library, 277–288 objects, 63-64 for the Stanford GraphBase, 285, 286 property_map class, 53, 156, 248-249 property_map.hpp, 56, 283 property_num class, 52 Property tags, 52, 155–156, 250, 285 property_traits class, 33, 56, 231, 278, 282 Prototypes, 43, 122, 164, 262 PROTOTYPES change file, 122, 262 push dispatch function, 234 push function, 234, 290 push-relabel algorithm, 105 push relabel flow function, 206 push_relabel_max_flow function, 209-212 push relabel max flow.hpp, 209 put function, 279, 280, 286 put_get_helper class, 286

R

RandomAccessIterator, 25-26, 27, 36 rank_map, 190 reachable_from_head vector, 72 read graphviz function, 84, 91 ReadWritePropertyMap, 103 Real model, 21 reference type, 29 Refinement, of concepts, 28, 138 remove_edge function, 152-153, 226, 228 remove edge if function, 227 remove_in_edge_if function, 155 remove_out_edge_if function, 154 remove_vertex function, 152, 225, 226 res_cap object, 108 Residual capacity, of edges, 105, 206 residual_capacity_map parameter, 206 rev_edge object, 108 reverse_edge_map parameter, 207 reverse graph adaptor, 13, 72 reverse_graph.cpp, 252-253 reverse_graph.hpp, 253 RIP (Routing Information Protocol), 76 roget_components.cpp, 262 root_map, 199 root vertex parameter, 194 Routers. See also Routing basic description of, 76 shortest-path problems and, 76–77 Routing. See also Routers distance vector, 77-81 link-state, 81–88 protocols, 76 tables, 76, 85-88 Routing Information Protocol (RIP), 76 Run-time behavior, testing, 126 dispatch, of virtual functions, 22

S

safe sort function, 35 Saturated edges, 105 Scherer, Andreas, 262 Self-loops, 4 Semicolon (;), 73 setS argument, 65, 127, 132 SGB (Stanford GraphBase), 119, 120, 122–123, 262 - 266sgb_vertex_id_map class, 122 sgb vertex name map class, 286 SGI STL Web site, 28 Shortest path. See also Paths; Shortest-path problems distance, 61 tree, 75 use of the term, 61, 63 weight, 75 Shortest-path problems. See also Paths; Shortest path basic description of, 61, 63, 75-88, 177-189 definitions, 75-76 Internet routing and, 76–77 single-pair, 75 single-source, 75 sink vertices, 105 "Six Degrees of Kevin Bacon" game, 62-67 size function, 290 sort function, 10, 34 Source code files bellman_ford.cpp, 185-186 container_gen.cpp, 232, 233 dijkstra.cpp, 227 graph.cpp, 126 incremental components.cpp, 202 interior_property_map.cpp, 232 kruskal.cpp, 191

miles_span.cpp, 262 prim.cpp, 195 reverse_graph.cpp, 252-253 roget_components.cpp, 262 source function, 124, 125 complexity guarantees and, 145 EdgeListGraph concept and, 147 undirected graphs and, 138, 140 Source vertex, 105 Spanning tree basic description of, 89 minimum-, problems, 89-96, 189-195 spanning_tree_edges iterator, 189 Specialization, partial, providing array traits for pointer types with, 32–33 Stack, basic description of, 19 Stanford GraphBase (SGB), 119, 120, 122-123, 262-266 stanford_graph.hpp, 13, 122, 243 std::accumulate function, 54 std::advance function, 33 std::back inserter function, 92 std::back insert iterator, 71 std::deque, 16 std::for_each function, 29 std::insert iterator, 109 std::istream iterator, 43 std::iterator_traits class, 29 std::list, 17, 23, 27, 127 std::map, 64, 84, 103 std::pair, 7-9, 43 AdjacencyGraph concept and, 146 AdjacencyMatrix concept and, 148–149 IncidenceGraph concept and, 144 interfacing with other graph libraries and, 126 shortest-path problems and, 78 std::set, 107, 109, 127, 132

std::sort, 36, 37 std::vector, 17, 27, 65, 85, 92, 107, 127 Stepanov, A. A., 25 STL (Standard Template Library) generic programming and, 22–23, 25–27 graph class comparisons and, 127 Graph concept and, 142 graph traversal and, 7 iterator_traits class, 29, 33, 278 traits class and, 33 visitors and, 10-11, 49 Web site, 28 strong_components function basic description of, 102–103, 198–201 named parameters, 199-200 parameters, 199 time complexity and, 200 strong_components.hpp, 198 Succ_Adj_Edge function, 125 Successors, number of, 117–118 sum function, 20, 21, 30, 31, 32

Т

Tags. See also Tags (listed by name) basic description of, 155 dispatching, 33–34 property, 52, 155–156, 250, 285
Tags (listed by name). See also Tags edge_length_t tag, 266 edge_weight_t tag, 52, 92 edge_xxx_t tag, 52 vertex_index_t tag, 223, 230, 266, 283 vertex_name_t tag, 52
target function, 124, 125 complexity guarantees and, 145 EdgeListGraph concept and, 149 undirected graphs and, 138, 140
TCP (Transmission Control Protocol), 76 Telephone lines, computing the best layout for, 90-96 Template(s). See also STL (Standard Template Library); Template parameters concept checking and, 34-37 for LEDA graphs, 266–272 polymorphism and, 21 size of, 23 specialization, 31 third-party, 32 traits class, 31-32 virtual functions and, comparison of, 22 visitors and, 50-51 Template parameters. See also Templates adjacency_list_traits class, 246 adjacency_matrix class, 238 adjacency_matrix_traits class, 247 disjoint_sets class, 294 edge list class, 251 filtered_graph class, 259 graph_property_iter_range class, 298 iterator_property_map class, 284 LEDA Graph class, 268 mutable_queue adaptor, 292 property class, 250 property_map class, 249 property_traits class, 283 reverse graph class, 253 Testing with graph class comparisons, 127 - 132run-time behavior, 126 tie function, 8, 296 Time. See also Time complexity compilation, 22, 54-55, 57-59 discover, of vertices, 44-45 finish, 44-45

Time complexity adjacency list class and, 225-227 basic description of, 165 bellman_ford_shortest_paths function and, 185 breadth_first_search function and, 167 connected_components function and, 196 depth_first_search function and, 172 depth_first_visit function and, 176 dijkstra_shortest_paths function and, 181 disjoint sets class and, 295 incremental components and, 202, 203, 204 johnson_all_pairs_shortest_paths function and, 188 kruskal_minimum_spanning_tree function and, 191 prim_minimum_spanning tree function and, 194 strong_components function and, 200 topological sort function and, 177 Timestamps, 116–117 Timing runs, 127–128 Tokens, 63 top function, 290 Topological sort. See also topological_sort function adjacency_list class and, 17 basic description of, 13-18, 176-177 via depth-first search, 18, 44-46 generic, 55-57 used with a vector of lists, 14–17 topological_sort function, 13, 14-18, 57. See also Topological sort basic description of, 176–177 depth first search algorithm and, 18, 44-46 interfacing with other graph libraries and, 119, 120-123

named parameters, 176–177 parameters, 176 topo_order, 46, 57-58, 121 topo_sort_dfs function, 46-47 topo sort function, 47, 49, 51, 55-58 topo_sort_visitor parameter, 18 topo_visitor, 56 tracert, 76 Traits class associated types and, 30-34 definition of, 31-32 Graph concept and, 142 internal properties and, 52 most well-known use of, 33 partial specialization and, 32-33 requirements and, 28 Transmission Control Protocol (TCP), 76 Traversal, 7-8, 24, 124, 137-150, 244 traversal_category function, 34, 124, 244 tree edge function, 50, 160 Tree edges, 61, 67 TrivialIterator, 28-29 Typedefs, nested in classes, 30-31

U

Undirected graphs, 4, 138–142 undirectedS argument, 63 Union by rank heuristics, 189–190, 293 union_set function, 294 User-defined objects, 49

V

Valid expressions, 28, 29 valid_position function, 114–115 value_type, 28, 29, 31 ValueType parameter, 251 vecS argument, 17, 53, 127 vector_as_graph.hpp, 15, 17, 273 Vector model, 21 vertex descriptor, 142, 144, 148 vertex function, 225 VertexGraph interface, 107 vertex index map function, 16, 187, 191, 194, 196, 199, 207 vertex_index_t tag, 223, 230, 266, 283 VertexList, 12, 127, 132, 134, 224-225, 230 VertexListGraph concept, 92, 124, 143, 147 vertex_list_graph_tag function, 124 VertexMutableGraph concept, 152 VertexMutablePropertyGraph concept, 156-157 vertex_name_t tag, 52 VertexProperties parameter, 12, 52 vertex_property, 157 VertexProperty parameter, 238 VertextMutableGraph, 124 vertex.t, 57 vertex xxx t, 52 Vertices accessing, 46-47 adding, 128–129 adjacent, 4, 46-47 basic description of, 3 clearing, 130, 134, 153 discover time of, 44-45 finish time of, 44–45 marking, using external properties, 46

names of, storing, 52, 53 sets of, basic description of, 3 source, 4 target, 4 traversing, 47-48, 130-132 vertices function, 8, 47, 296 vertices_size_type function, 124 Virtual functions compile-time dispatch of, 22 run-time dispatch of, 22 size of, 22 Visitor(s), 10–11, 50 basic description of, 49-52 concepts, 158-160 visitor function, 11, 66 visitor parameter, 162, 184 Visual C++ (Microsoft), 41, 127, 128

W

Warnsdorff's heuristic, 113, 117–118 Web page(s) connected components and, 97, 102–104 links, connected components and, 102–104 well-designed, 97 weight_map parameter, 79, 184, 187, 194 WeightMap type, 190 white_color constant, 56 who_owes_who function, 231