第36卷第3期 中国科学技术大学学报 Vol. 36, No. 3
2006年3月 JOURNAL OF UNIVERSITY OF SCIENCE AND TECHNOLOGY OF CHINA Mar. 2006

Article ID:0253-2778(2006)03-0249-05

On restricted edge connectivity and extra edge connectivity of hypercubes and folded hypercubes*

ZHU Qiang[1,2], XU Jun-ming[1]

(1. *Department of Mathematics, University of Science and Technology of China, Hefei* 230026, *China*;
2. *Department of Mathematics, Xidian University, Xi'an* 710071, *China*)

Abstract: The 2-extra edge connectivity of the hypercubes and the 1-extra edge connectivity and restricted edge connectivity of the folded hypercubes are determined.

Key words: hypercube; restricted edge connectivity; extra edge connectivity

CLC number: O157.5; TP302.1 **Document code**: A

AMS Subject Classification(2000): Primary 05C40; Secondary 90B10

立方体和折叠立方体的限制边连通度和超边连通度

朱 强[1,2],徐俊明[1]

(1.中国科学技术大学数学系,安徽合肥 230026;2.西安电子科技大学数学系,陕西西安 710071)

摘要:确定了立方体的2-超边连通度和折叠立方体的1-超边连通度和限制边连通度.

关键词:立方体;限制边连通度;超边连通度

0 Introduction

For all the terminologies and notations not defined here, we follow ref. [1]. For a graph $G=(V, E)$ and $S \subset V(G)$ or $S \subset G$, we use $E_G(S)$ to denote the set of neighboring edges of S in G, that is, $E_G(S)=\{xy: y \in V(G-S), xy \in E(G)$ for some $x \in S\}$. In this paper, we use graph and interconnection networks, nodes and vertices, links and edges interchangeably.

Edge connectivity is an important parameter to measure the fault tolerant ability of interconnection networks, however, in many cases this parameter greatly underestimates this ability. Since in many practical applications it can be safely assumed that any set of faults in some networks can not contain all links which are directly connected to some processor. For these networks, the classical edge connectivity may not be accurate measures of network reliability[2~4]. To compensate for this shortcoming, ref. [2] proposed the concept of the restricted edge connectivity $\lambda^1(G)$ of G. A subset $S \subset E(G)$ is called a restricted edge set, if S does not contain the neighboring edge set of any vertex as

* **Received**: 2004-02-20; **Revised**: 2005-10-25
Foundation item: Supported by NNSF of China (10271114, 10301031).
Biography: ZHU Qiang, male, born in 1976, PhD. Research field: parallel computing, interconnection networks, graphy theory, combinatorial optimization. E-mail: qiangzhu@ustc.edu
Corresponding author: XU Jun-ming, Prof. E-mail: xujm@ustc.edu.cn

its subset; and a restricted edge set S is called a restricted edge cut if $G-S$ is disconnected. If there exists a restricted edge cut S in G, then the restricted edge connectivity $\lambda^1(G) = \min\{|S| : S$ is a restricted edge cut of $G\}$. Ref. [3] determined that $\lambda^1(Q_n) = 2n-2$ for $n \geqslant 3$, where Q_n is an n-dimensional hypercube. As far as we know, the restricted edge connectivity of the folded hypercubes has not been determined.

Ref. [5] has generalized the concept of the restricted edge connectivity to the h-restricted edge connectivity. A subset $S \subset E(G)$ is called an h-restricted edge cut if $G-S$ is disconnected and every remaining component has the minimum degree of vertex at least h. If there exists an h-restricted edge cut S in G, then the h-restricted edge connectivity $\lambda^h(G) = \min\{|S| : S$ is an h-restricted edge cut of $G\}$. Ref. [5] has determined that the h-restricted connectivity of Q_n is $\kappa^h(Q_n) = (n-h)2^h$ for $0 \leqslant h \leqslant \left\lfloor \frac{n}{2} \right\rfloor$. But they did not determine the h-restricted edge connectivity of Q_n.

Ref. [6] has defined $\lambda(G; \mathcal{P}_h)$ for a given non-negative integer h as the minimum cardinality of a set of edges, if any, whose deletion disconnects G and every remaining component has more than h vertices. And they called this type of conditional edge connectivity as h-extra edge connectivity of G, denoted by $\lambda_h(G)$. To be exact, an edge-cut S of G is called a $\mathcal{P}_h$ edge-cut if every component of $G-S$ has more than h vertices. If there exists a $\mathcal{P}_h$ edge-cut in G, then the h-extra connectivity $\lambda_h(G) = \min\{|S| : S$ is a $\mathcal{P}_h$ edge-cut of $G\}$. By definition, a graph G having the property $\mathcal{P}_0$ implies that every component of G contains at least one vertex. Thus, $\lambda_0(G) = \lambda(G)$ if G is not a complete graph.

Since for any graph G and any edge subset F, the condition that there is no isolated vertex in $G-F$ and the condition that F does not contain the neighboring edge set of any vertex is equivalent, so a $\mathcal{P}_1$ edge cut of G is also a restricted edge cut of G, thus for any graph G, if $\lambda_1(G)$ or $\lambda_G(1)$ exists, then:

Lemma 0.1 $\lambda^1(G) = \lambda_1(G)$.

We are, in this paper, interested in the hypercube Q_n and the folded hypercube FQ_n, which have been widely used in the design and analysis of interconnection networks. As we have already known, $\lambda(Q_n) = n$. Ref. [7] determined $\lambda(FQ_n) = n+1$. Ref. [3] has determined $\lambda^1(Q_n) = 2n-2$ for $n \geqslant 3$. Thus by Lemma 0.1, $\lambda_1(Q_n) = 2n-2$ for $n \geqslant 3$. In this paper, we will determine $\lambda_2(Q_n) = 3n-4$ for $n \geqslant 4$, $\lambda_1(FQ_n) = \lambda^1(FQ_n) = 2n$ for $n \geqslant 4$. The result $\lambda_2(Q_n) = 3n-4$ means that the hypercube can tolerate $3n-4$ link failures without being disconnected provided that all the neighboring edges of any subtree with order not more than 2 can't fail at the same time. This result greatly improves the fault tolerant ability of n-cube theoretically.

1 Results on hypercube

An n-dimensional hypercube (i. e., n-cube) Q_n can be modelled as a graph $G_n(V, E)$, with $|V| = N = 2^n$, and $|E| = n2^{n-1}$. Each node represents a processor and each edge represents a link between a pair of processors. Nodes are assigned binary numbers from 0 to $2^n - 1$ such that labels of any two neighbors differ only in one bit position. Links are also labelled from 0 to $n-1$ such that any link labelled i connects two nodes whose labels differ in the i^{th} bit. Since the rightmost bit position is the 0^{th} position, so all the links labelled 0 are cross edges. The neighbors of a node u are called bordering nodes of u. The bordering node of u across dimension i is denoted by u_i. The bordering node of u_i across dimension j is denoted by u_{ij} (Thus $u_{ii} = u$). The neighboring links of a node u are called bordering links of u, the bordering links of u across dimension i is denoted by $e_i(u)$, since $e_0(u)$ is a cross edge, we also denote $e_0(u)$ by $e_c(u)$.

Theorem 1.1 $\lambda_2(Q_n) = 3n-4, n \geqslant 4$.

Proof 1) Find a path $P_2 = (u \to v \to w)$ of length 2 in Q_n, it is easy to see that $|E_{Q_n}(P_2)| =$

$(n-1)+(n-2)+(n-1)=3n-4$. Since $\lambda(Q_n-\{u, v, w\}) \geqslant \kappa(Q_n-\{u, v, w\}) \geqslant n-3>0$ (when $n \geqslant 4$), so Q_n-P_2 is connected, so $Q_n-E_{Q_n}(P_2)$ is disconnected and contains no isolated vertex or isolated edge, $\lambda_2(Q_n) \leqslant |E_{Q_n}(P_2)|=3n-4$.

2) Suppose $A \subset E(Q_n)$, $|A|=3n-5$, and there is no isolated vertex or isolated edge in Q_n-A. In the following, we will prove that Q_n-A is connected. Following ref. [3], we express Q_n as $Q_n=L \oplus R$, where L and R are the two $(n-1)$-subcubes of Q_n induced by the vertices with the leftmost coordinate 0 and 1, respectively, that is, all the vertices in L are of the form $0*\cdots*$ and all the vertices in R are of the form $1*\cdots*$. Let $A_L=A \cap L$, $A_R=A \cap R$. Since $L \cap R=\varnothing$, so $|A_L|+|A_R| \leqslant |A|=3n-5 \leqslant 4n-9$ (when $n \geqslant 4$), then either $|A_L| \leqslant 2n-5$ or $|A_R| \leqslant 2n-5$. Without loss of generality, we suppose that $|A_R| \leqslant 2n-5$.

In the following, we will prove that the vertices in $R-A_R$ is connected to each other in Q_n-A.

Case 1 If there is no isolated vertex in $R-A_R$, then by ref. [3], $\lambda'(R)=\lambda'(Q_{n-1})=2n-4>2n-5=|A_R|$, so $R-A_R$ is a connected graph, we are done.

Case 2 If there exists an isolated vertex u^R in $R-A_R$, then $\lambda(R-u^R) \geqslant \kappa(R-u^R) \geqslant \kappa(R)-1=n-2$, and $|A_R|-|E(u^R:R)| \leqslant 2n-5-(n-1)=n-4$. So at most $n-4 (\leqslant \lambda(R-u^R))$ edges of $R-u^R$ may be faulty (may be in A_R), so the subgraph $G=(R-u^R)-A_R=(R-A_R)-u^R$ is connected. In the following we will prove that u^R is connected to the subgraph G in Q_n-A. Since there is no isolated vertex in Q_n-A, the cross edge $e_c(u^R)=(u^R, u^L) \notin A$. For $1 \leqslant i \leqslant n-1$, if there exists an i such that both $e_i(u^L) \notin A$ and $e_c(u_i^L) \notin A$ (u_i^L mean the neighboring vertex of u^L which differs from u^L in the i^{th} bit), then u^R can be connected to G by the path: $u^R \xrightarrow{e_c(u^R)} u^L \xrightarrow{e_i(u^L)} u_i^L \xrightarrow{e_c(u_i^L)} G=(R-u^R)-A_R$, we are done. So we may suppose that for each i at least one of $e_i(u^L)$ and $e_c(u_i^L)$ is in A. Let $B=\{e_i(u^L), e_c(u_i^L) \mid i=1, 2, \cdots, n-1\} \cap A$, then $|B| \geqslant n-1$. Since there is no isolated edge in Q_n-A, there exists a j such that $e_j(u^L)=(u^L, v^L) \notin A_L$. let $C=\{e_i(v^L), e_c(v_i^L) \mid i=1, 2, \cdots, n-1$ but $i \neq j\}$ it is obvious that $E(u^R:R)$ (the neighboring vertex set of u^R in R), B and C are disjoint. So $|C \cap A| \leqslant |A|-|A(u^R:R)|-|B| \leqslant n-3$. Since there are $n-2$ pair of edges $(e_i(v^L), e_c(v_i^L))$ in C, so there exists an i_1 such that neither $e_{i_1}(v^L)$ nor $e_c(v_{i_1}{}^L)$ belonging to A, u_R can be connected to the connected subgraph $(R-A_R)-u_R$ through the edges (u_R, u_L), (u_L, v_L), $e_{i_1}(v^L)$ and $e_c(v_{i_1}{}^L)$, thus completing our proof that the vertices in $R-A_R$ is connected to each other in Q_n-A.

In the following paragraph, we will prove that any vertex of $L-A_L$ is connected to the subgraph $R-A_R$.

Suppose that x^l is any vertex in $L-A_L$, if $e_c(x^l) \notin A$, then we are done. So we suppose that $e_c(x^l) \in A$. If there exists an $i \in \{1, 2, \cdots, n-1\}$ such that both $e_i(x^l) \notin A$ and $e_c(x_i{}^l) \notin A$, then we are done. So we suppose that at least one edge from each of the above $n-1$ pairs belongs to A. Let $B'=\{e_i(x^l), e_c(x_i{}^l) \mid i=1, 2, \cdots, n-1\} \cap A$, then $|B'| \geqslant n-1$. Since there is no isolated vertex in Q_n-A, there exists a j such that $e_j(x^l) \notin A$. Suppose that $e_j(x^l)=(x^l, y^l)$. For all the $i \in \{1, 2, \cdots, n-1\}$ and $i \neq j$, if both $e_i(y^l)$ and $e_c(y_i^l)$ do not belong to A, then we are done. So we suppose that at least one edge from each of the above $n-2$ pairs belongs to A. Let $C'=\{e_i(y^l), e_c(y_i{}^l) \mid i=1, 2, \cdots, j-1, j+1, \cdots, n-1\} \cap A$, then $|C'| \geqslant n-2$. Since there is no isolated edge in Q_n-A, so $D=E(e_j(x^l))-A_R \neq \varnothing$ (where $E(e_j(x^l))$ denotes the neighboring edge set of the edge $e_j(x^l)$). Suppose that z^l is an end-vertex of an element of D, and $z^l \notin \{x^l, y^l\}$. Since there is no triangle in Q_n, the vertex z^l can be adjacent to just one of the two vertices x^l, y^l. Let $E'=\{e_i(z^l), e_c(z_i{}^l) \mid i \in \{1, 2, \cdots, n-1\}$ and $e_i(z^l)$ is not incident to x^l or $y^l\}$. Since $\{e_c(x^l)\}$, B', C', E' are disjoint to each other, so $|E' \cap A| \leqslant 3n-5-(1+n-1+n-2)$

$= n-3$. Since there are $n-2$ pairs of edges in E', so there exists a pair of edges (z^l, z_k^l) and $e_c(z_k^l)$, neither of which belonging to A, so z^l can be connected to the subgraph $R-A_R$. Since x^l can be connected to z^l, our proof is complete.

2 Results on folded hypercube

An n-dimensional folded hypercube FQ_n is basically a standard hypercube augmented with some extra links between nodes. There are 2^{n-1} links between all pairs of complementary nodes. Two nodes in a hypercube are said to be complementary if the exclusive-OR of their addresses gets all 1s.

Theorem 2.1 $\lambda^1(FQ_n)=\lambda_1(FQ_n)=2n, n\geqslant 4$.

Proof 1) Let $e=(x, y)$ be an edge of FQ_n, then it is obvious that $FQ_n - E_{FQ_n}(e)$ is disconnected since e is a component of it. Since $\kappa(FQ_n-\{x, y\})\geqslant\kappa(FQ_n)-2=n-1>0$ (when $n\geqslant 3$), so there is no isolated vertex in $FQ_n - E_{FQ_n}(e)$, $\lambda_1(FQ_n)\leqslant 2n$.

2) Suppose $F\subset E(FQ_n)$, $|F|\leqslant 2n-1$, and there is no isolated vertex in $FQ_n - F$. In the following we will prove that $FQ_n - F$ is connected. Since at least $2n$ edges are to be removed to get an isolated edge in FQ_n, there is no isolated edge in FQ_n-F.

First we define a map ϕ between $V(Q_{n+1})$ and $V(FQ_n)$, We use the symbol u to represent an n-bit binary string, $\bar{u}$ to represent u's complementary n-bit binary string. We define ϕ as follows:

$$\phi(0u)=u,\ \phi(1u)=\bar{u}.$$

It is easy to verify that the map ϕ induces another map ρ from $E(Q_{n+1})$ to $E(FQ_n)$:

$$\rho(e)=(\phi(x),\phi(y))\in E(FQ_n)$$

where $e=(x, y)$ is an edge of Q_{n+1}.

So ϕ is a graph homomorphism between Q_{n+1} and FQ_n. For any edge $e=(u, v)\in E(FQ_n)$, if e is not a complementary edge (the two end-vertices of a complementary edge are complementary n-bit binary strings), then $\rho^-(e)=\{(0u, 0v), (1\bar{u}, 1\bar{v})\}$ if e is a complementary edge, then $u=\bar{v}$, so $\rho^-(e)=\{(0u, 1\bar{v}), (0v, 1\bar{u})\}=\{(0u, 1u), (0v, 1v)\}$ is a set of 2 cross edges in Q_{n+1}. For any edge set $F\subset E(FQ_n)$, we define $\rho^-(F)=\bigcup_{e\in F}\rho^-(e)$.

We express Q_{n+1} as $Q_{n+1}=L\oplus R$, where L and R are the two n-subcubes of Q_{n+1} induced by the vertices with the leftmost coordinate 0 and 1 respectively, and we define $F_l=\rho^-(F)\cap E(L)$ and $F_r=\rho^-(F)\cap E(R)$. We define ψ as a map from $V(L)$ to $V(R)$: $\psi(0u)=1\bar{u}$. It is easy to verify that ψ is a graph isomorphism between $L-F_l$ and $R-F_r$. Since $F_l\cap F_r=\phi$, $|F_l|+|F_r|\leqslant|\rho^-(F)|=2(2n-1)$, and $|F_l|=|F_r|$, so $|F_l|=|F_r|\leqslant 2n-1$.

For any two vertex u, v in FQ_n-F, if there is a path P between $0u$ and $0v$ in $Q_{n+1}-\rho^-(F)$, then it is easy to see that $\rho(P)$ is a path between u and v in FQ_n-F, so if $Q_{n+1}-\rho^-(F)$ is connected, then FQ_n-F is connected, too. In the following we will prove that $Q_{n+1}-\rho^-(F)$ is connected.

Consider the graph $L-F_l$ and suppose that $e=(u, v)$ is any edge in L. Since $\lambda(L-\{u, v\})\geqslant\kappa(L-\{u, v\})\geqslant n-2>1\geqslant|F_l-E_L(e)|$ (when $n\geqslant 4$). So if there exists an isolated edge $e=(u, v)$ in $L-F_l$, then $L-F_l$ has just 2 components.

Since $|F_l|\leqslant 2n-1$, there are at most two isolated vertices in $L-F_l$. And if there are two isolated vertices in $L-F_l$, the 2 isolated vertices must be adjacent. In the following we will consider 4 cases and prove that $Q_{n+1}-\rho^-(F)$ is connected in all these cases. A) There are 2 isolated vertices in $L-F_l$; B) There is just 1 isolated vertex in $L-F_l$; C) There is an isolated edge in $L-F_l$, and D) There are no isolated vertex or isolated edge in $L-F_l$.

A) If there are 2 isolated vertices $0u$ and $0v$, from the above we know that $(0u, 0v)\in E(L)$ and $F_l=E_L(0u)\cup E_L(0v)$, and since $\kappa(L-\{0u, 0v\})\geqslant n-2>0$ (when $n\geqslant 3$), so $L-\{0u, 0v\}=L-F_l-\{0u, 0v\}$ is connected. Since $L-F_l$ is isomorphism to $R-F_r$ under the map ψ, so the two isolated vertices in $R-F_r$ are $\psi(0u)=1\bar{u}$ and $\psi(0v)=1\bar{v}$ and $R-\{1\bar{u}, 1\bar{v}\}$ is connected. Since $|F_l|=2n-1$, there is no element of $\rho^-(F)$ which is a cross edge between L and R. $0u$ and $0v$ are connected to the connected subgraph $R-\{1\bar{u}, 1\bar{v}\}$ via $(0u, 1u)$ and $(0v, 1v)$ respectively. The 2 isolated vertices in $R-F_r$ are connected to the connected subgraph $L-\{u, v\}$ of $L-F_l$ via $(0\bar{u}, 1\bar{u})$ and $(0\bar{v}, 1\bar{v})$ respectively. And since the cross edges between $L-\{u, v\}$ and $R-\{1\bar{u}, 1\bar{v}\}$ do not belong to $\rho^-(F)$, so the two connected subgraph are connected to each other, thus completing our proof, $Q_{n+1}-\rho^-(F)$ is connected.

B) If there is just 1 isolated vertex $0u$ in $L-F_l$, then it's easy to know that there is no isolated edge in $L-F_l$. In the following we will prove that $(L-F_l)-0u$ is connected. For otherwise, we define $F'=F_l-e_1(0u)$, in the graph $L-F'$, $0u$ is connected to only one of its components *via* the edge $e_1(0u)$. $L-F'$ is disconnected since $(L-F_l)-0u$ is disconnected. And there is no isolated vertex or isolated edge in $L-F'$ since $0u$ is the only isolated vertex in $L-F_l$ and there is no isolated edge in $L-F_l$. Since $|F'|=|F_l|-1\leqslant 2n-2<\lambda_2(Q_n)=3n-4$ (when $n\geqslant 3$), thus we have obtained a contradiction, which means that $(L-F_l)-0u$ is connected. Since $L-F_l$ is isomorphic to $R-F_r$ under the map ψ, so $\psi(0u)=1\bar{u}$ is an isolated vertex in $R-F_r$ and $R-F_r-1\bar{u}$ is connected. Since there are no isolated vertices in $Q_{n+1}-\rho^-(F)$, $0u$ is connected to the connected subgraph $R-F_r-1\bar{u}$ and $1\bar{u}$ is connected to the connected subgraph $L-F_l-0u$, since there are 2^n-2 cross edges between $L-F_l-0u$ and $R-F_r-1\bar{u}$, at most $2(2n-1-n)=2n-2$ of them may be in $\rho^-(F)$. Since $2^n-2>2n-2$ (when $n\geqslant 3$), $L-F_l-0u$ is connected to $R-F_r-1\bar{u}$. That is, $Q_{n+1}-\rho^-(F)$ is connected.

C) If there exists an isolated edge $e=(0u, 0v)$ in $L-F_l$, since $|F_l|\geqslant |E_L(e)|=2n-2$, at most 2 cross edges belong to $\rho^-(F)$. Consider the graph $L-F_l-\{0u, 0v\}$, since $\lambda(L-\{0u, 0v\})\geqslant \kappa(L-\{0u, 0v\})\geqslant n-2>1$ (when $n>3$), so $L-F_l-\{0u, 0v\}$ is connected. By the isomorphism of $L-F_l$ and $R-F_r$, $\bar{e}=(1\bar{u}, 1\bar{v})$ is an isolated edge in $R-F_r$ and $R-F_r-\{1\bar{u}, 1\bar{v}\}$ is a connected subgraph. According to the definition of the map ρ, at least one of $(0u, 1u)$ and $(0v, 1v)$ does not belong to $\rho^-(F)$ (or (u,v) will be an isolated edge in $F_{Q_n}-F$, a contradiction). Thus e is connected to $R-F_r-\{1\bar{u}, 1\bar{v}\}$. Similarly the edge $\bar{e}$ is connected to $L-F_l-\{0u, 0v\}$. There are 2^n-4 cross edges between $L-F_l-\{0u, 0v\}$ and $R-F_r-\{1\bar{u}, 1\bar{v}\}$. At most 2 of them can be in $\rho^-(F)$, so $L-F_l-\{0u, 0v\}$ and $R-F_r-\{1\bar{u}, 1\bar{v}\}$ is connected to each other. Thus we have proved that $Q_{n+1}-\rho^-(F)$ is connected in this case.

D) If there is no isolated vertex or isolated edge in $L-F_l$, since $|F_l|\leqslant 2n-1<3n-4$ (when $n>3$), so $L-F_l$ is connected. By the isomorphism of $L-F_l$ and $R-F_r$, $R-F_r$ is connected too. There are 2^n cross edges between $L-F_l$ and $R-F_r$, at most $2(2n-1)<2^n$ (when $n>3$), so $L-F_l$ is connected to $R-F_r$. Thus we have proved in this case that $Q_{n+1}-\rho^-(F)$ is connected.

The above 4 cases complete our proof that $Q_{n+1}-\rho^-(F)$ is connected, so FQ_n-F is connected, so $\lambda_1(FQ_n)\geqslant 2n$.

By 1) and 2), we have proved that $\lambda_1(FQ_n)=\lambda_1(FQ_n)=2n$. □

References

[1] Bondy J A, Murty U S R. Graph Theory with Applications[M]. New York: Elsevier, 1976.

[2] Esfahanian A H, Hakimi S L. On computer a conditional edge connectivity of graph[J]. Information Processing Letters, 1988,27(4):195-199.

[3] Esfahanian A H. Generalized measures of fault tolerance with applications to N-cube networks[J]. IEEE Transactions on Computers,1989,38(11):1 586-1 591.

[4] Harary F. Conditional connectivity[J]. Networks, 1983,13:346-357.

[5] Latifi S, Hegde M, Naraghi-Pour M. Conditional connectivity measures for large multiprocessor systems [J]. IEEE Transactions on Computers, 1994, 43(2): 218-221.

[6] Fàbrega J, Fiol M A. Extraconnectivity of graphs with large girth[J]. Discrete Math., 1994, 127:163-170.

[7] El-Amawy A, Latifi S. Properties and performance of folded hypercubes[J]. IEEE Transactions on Parallel and Distributed Systems, 1991, 2(1):31-42.

[8] XU Jun-ming. Topological Structure and Analysis of Interconnection Networks [M]. Dordrecht/Boston/London: Kluwer Academic Publishers, 2001.