Notation and Conventions

We use standard notation in this book. The comments here are mostly for clarification. You might want to skip this section and return to it only as needed.

Sets

We write $[N] := \{1, 2, \ldots, N\}$. Also $\mathbb{N} := \{1, 2, \ldots\}$.

Given a finite set S and a positive integer r, we write $\binom{S}{r}$ for the set of r-element subsets of S.

If S is a finite set and f is a function on S, we use the expectation notation $\mathbb{E}_{x \in S} f(x)$, or more simply $\mathbb{E}_x f(x)$ (or even $\mathbb{E} f$ if there is no confusion) to mean the average $|S|^{-1} \sum_{x \in S} f(x)$. We also use the symbol $\mathbb{E}$ for its usual meaning as the expectation for some random variable.

A ***k*-term arithmetic progression** (abbreviated ***k*-AP**) in an abelian group is a sequence of the form

$$a, a + d, a + 2d, \ldots, a + (k - 1)d.$$

Here d is called the common difference. The progression is called ***nontrivial*** if $d \neq 0$, and ***trivial*** if $d = 0$. When we say that a set A contains a k-AP, we mean that it contains a nontrivial k-AP. Likewise, when we say that A is ***k*-AP-free**, we mean that it contains no nontrivial k-APs.

Graphs

We write a graph as $G = (V, E)$, where V is a finite set of vertices, and E is the set of edges. Each edge is an unordered pair of distinct vertices. Formally, $E \subseteq \binom{V}{2}$.

Given a graph G, we write $V(G)$ for the set of vertices, and $E(G)$ for the set of edges, and denote their cardinalities by $v(G) := |V(G)|$ and $e(G) := |E(G)|$.

In a graph G, the ***neighborhood*** of a vertex x, denoted $N_G(x)$ (or simply $N(x)$ if there is no confusion), is the set of vertices y such that xy is an edge. The ***degree*** of x is the number of neighbors of x, denoted $\deg_G(x) := |N_G(x)|$ (or simply written as $\deg(x)$).

Given a graph G, for each $A \subseteq V(G)$, we write $e(A)$ to denote the number of edges with both endpoints in A. Given $A, B \subseteq V(G)$ (not necessarily disjoint), we write

$$e(A, B) := |\{(a, b) \in A \times B : ab \in E(G)\}|.$$

Note that when A and B are disjoint, $e(A, B)$ is the number of the edges between A and B. On the other hand, $e(A, A) = 2e(A)$ as each edge within A is counted twice.

Here are some standard graphs:

- ***K_r*** is the complete graph on r vertices, also known as an ***r-clique***;
- ***$K_{s,t}$*** is the complete bipartite graph with s vertices in one vertex part and t vertices in the other vertex part;
- ***$K_{r,s,t}$*** is a complete tripartite graph with vertex parts having sizes r, s, t respectively (e.g., $K_{1,1,1} = K_3$); and so on analogously for complete multipartite graphs with more parts;
- ***C_ℓ*** ($\ell \geq 3$) is a cycle with ℓ vertices and ℓ edges.

Some examples are shown below.

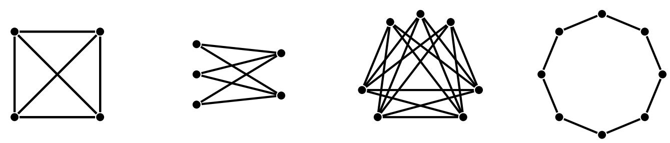

K_4 $K_{3,2}$ $K_{3,2,2}$ C_8

Given two graphs H and G, we say that H is a ***subgraph*** of G if one can delete some vertices and edges from G to obtain a graph isomorphic to H. A ***copy*** of H in G is a subgraph G that is isomorphic to H. A ***labeled copy*** of H in G is a subgraph of G isomorphic to H where we also specify the isomorphism from H. Equivalently, a labeled copy of H in G is an injective graph homomorphism from H to G. For example, if G has q copies of K_3, then G has $6q$ labeled copies of K_3.

We say that H is an ***induced subgraph*** of G if one can delete some vertices of G (when we delete a vertex, we also remove all edges incident to the vertex) to obtain H – note that in particular we are not allowed to remove additional edges other than those incident to a deleted vertex. If $S \subseteq V(G)$, we write ***$G[S]$*** to denote the subgraph of G induced by the vertex set S, that is, $G[S]$ is the subgraph with vertex set S and keeping all the edges from G among S.

As an example, the following graph contains the 4-cycle as an induced subgraph. It contains the 5-cycle as a subgraph but not as an induced subgraph.

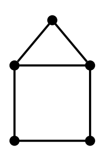

In this book, when we say ***H-free***, we always mean not containing H as a subgraph. On the other hand, we say ***induced H-free*** to mean not containing H as an induced subgraph.

Given two graphs F and G, a ***graph homomorphism*** is a map $\phi \colon V(F) \to V(G)$ (not necessarily injective) such that $\phi(u)\phi(v) \in E(G)$ whenever $uv \in E(F)$. In other words, ϕ is a map of vertices that sends edges to edges. A key difference between a copy of F in G and a graph homomorphism from F to G is that the latter does not have to be an injective map of vertices.

The ***chromatic number $\chi(G)$*** of a graph G is the smallest number of colors needed to color the vertices of G of so that no two adjacent vertices receive the same color (such a coloring is called a ***proper coloring***).

The ***adjacency matrix*** of a graph $G = (V, E)$ is a $v(G) \times v(G)$ matrix whose rows and columns both are indexed by V, and such that the entry indexed by $(u, v) \in V \times V$ is 1 if $uv \in E$ and 0 if $uv \notin E$.

An ***r-uniform hypergraph*** (also called ***r-graph*** for short) consists of a finite vertex set V along with an edge set $E \subseteq \binom{V}{r}$. Each edge of the r-graph is an r-element subset of vertices.

Asymptotics

We use the following standard asymptotic notation. Given nonnegative quantities f and g, in each of the following items, the various notations have the same meaning (as some parameter, usually n, tends to infinity):

- $\boldsymbol{f \lesssim g}$, $\boldsymbol{f = O(g)}$, $\boldsymbol{g = \Omega(f)}$, $f \le Cg$ for some constant $C > 0$
- $\boldsymbol{f = o(g)}$, $f/g \to 0$
- $\boldsymbol{f = \Theta(g)}$, $\boldsymbol{f \asymp g}$, $g \lesssim f \lesssim g$
- $\boldsymbol{f \sim g}$, $f = (1 + o(1))g$

Subscripts (e.g., $\boldsymbol{O_s(\ )}$, $\boldsymbol{\lesssim_s}$) are used to emphasize that the hidden constants may depend on the subscripted parameters. For example, $f(s, x) \lesssim_s g(s, x)$ means that for every s there is some constant C_s so that $f(s, x) \le C_s g(s, x)$ for all x.

We avoid using $\ll$ since this notation carries different meanings in different communities and by different authors. In analytic number theory, $f \ll g$ is standard for $f = O(g)$ (this is called Vinogradov notation). In combinatorics and probability, $f \ll g$ sometimes means $f = o(g)$, and sometimes means that f is sufficiently small depending on g.

When asymptotic notation is used in the hypothesis of a statement, it should be interpreted as being applied to a sequence rather than a single object. For example, given functions f and g, we write

$$\textit{if } f(G) = o(1), \textit{ then } g(G) = o(1)$$

to mean

$$\textit{if a sequence } G_n \textit{ satisfies } f(G_n) = o(1), \textit{ then } g(G_n) = o(1),$$

which is also equivalent to

$$\textit{for every } \varepsilon > 0 \textit{ there is some } \delta > 0 \textit{ such that, if } |f(G)| \le \delta, \textit{ then } |g(G)| \le \varepsilon.$$