Lecture 11

Hamiltonian graphs and the Bondy-Chvátal Theorem

This lecture introduces the notion of a Hamiltonian graph and proves a lovely theorem due to J. Adrian Bondy and Vašek Chvátal that says—in essence—that if a graph has lots of edges, then it must be Hamiltonian.

Reading:
The material in today's lecture comes from Section 1.4 of

> Dieter Jungnickel (2013), *Graphs, Networks and Algorithms*, 4th edition, (available online via `SpringerLink`),

and is essentially an expanded version of the proof of Jungnickel's Theorem 1.4.1.

11.1 Hamiltonian graphs

In the last lecture we characterised Eulerian graphs, which are those that have a closed trail that includes every *edge* exactly once. It's then natural to wonder about graphs that have closed trails that include every *vertex* exactly once. Somewhat surprisingly, these turn out to be much, much harder to characterise. To begin with, let's make some definitions that parallel those for Eulerian graphs:

Definition 11.1. *A* ***Hamiltonian path*** *in a graph $G(V, E)$ is a path that includes all of the graph's vertices.*

Definition 11.2. *A* ***Hamiltonian tour*** *or* ***Hamiltonian cycle*** *in a graph $G(V, E)$ is a cycle that includes every vertex.*

Definition 11.3. *A graph that contains a Hamiltonian tour is said to be a* ***Hamiltonian graph****. Note that this implies that Hamiltonian graphs have $|V| \geq 3$, as otherwise they would be unable to contain a cycle.*

Generally speaking, it's difficult to decide whether a graph is Hamiltonian—there are no known efficient algorithms. There are, however, some special cases that are easy: the cycle graphs C_n consist of nothing except one big Hamiltonian tour, and the complete graphs K_n with $n \geq 3$ obviously contain the Hamiltonian cycle

$$(v_1, v_2, \ldots, v_n, v_1)$$

obtained by numbering the vertices and visiting them in order. We'll spend most of the lecture proving results that say, more-or-less, that a graph with a lot of edges (where the point of the theorem is to make the sense of "a lot" precise) is Hamiltonian. Two of the simplest results of this kind are:

Theorem 11.4 (Dirac[1], 1952). *Let G be a graph with $n \geq 3$ vertices. If each vertex of G has $\deg(v) \geq n/2$, then G is Hamiltonian.*

Theorem 11.5 (Ore, 1960). *Let G be a graph with $n \geq 3$ vertices. If*

$$\deg(u) + \deg(v) \geq n$$

for every pair of non-adjacent vertices u and v, then G is Hamiltonian.

Dirac's theorem is a corollary of Ore's, but we will not prove either of these theorems directly. Instead, we'll obtain both as corollaries of a more general result, the Bondy-Chvátal Theorem. Before we can even formulate this mighty result, we need a somewhat involved new definition: the *closure* of a graph.

11.2 The closure a graph

Suppose G is a graph on n vertices. Then the ***closure*** of G, written $[G]$, is constructed by adding edges that connect pairs of non-adjacent vertices u and v for which

$$\deg(u) + \deg(v) \geq n. \tag{11.1}$$

One continues recursively, adding new edges according to (11.1) until all non-adjacent pairs u, v satisfy

$$\deg(u) + \deg(v) < n.$$

The graphs G and $[G]$ have the same vertex set—I'll call it V—but the edge set of $[G]$ may contain extra edges. In the next section I'll give an explicit algorithm that constructs the closure.

[1] This Dirac, Gabriel Andrew Dirac, was the adopted son of the Nobel prize winning theoretical physicist Paul A. M. Dirac, and the nephew of another Nobel prize winner, the physicist and mathematician Eugene Wigner. Wigner's sister Margit was visiting her brother in Princeton when she met Paul Dirac.

11.2.1 An algorithm to construct $[G]$

The algorithm below constructs a finite sequence of graphs

$$G = G_1(V, E_1),\ G_2(V, E_2),\ \ldots,\ G_K(V, E_K) = [G] \tag{11.2}$$

that all have the same vertex set V, but different edge sets

$$E = E_1,\ E_2,\ \ldots,\ E_K. \tag{11.3}$$

These edge sets form an increasing sequence in the sense that that $E_j \subset E_{j+1}$. In fact, E_{j+1} is produced by adding a single edge to E_j.

Algorithm 11.6 (Graph Closure)**.**
Given a graph $G(V,\ E)$ with vertex set $V = \{v_1,\ \ldots,\ v_n\}$, find $[G]$.

(1) Set an index j to one: $j \leftarrow 1$,
Also set E_1 to be the edge set of the original graph,
$E_1 \leftarrow E$.

(2) Given E_j, construct E_{j+1}, which contains, at most, one more edge than E_j. Begin by setting $E_{j+1} \leftarrow E_j$, so that E_{j+1} automatically includes every edge in E_j. Now work through every possible edge in the graph. For each one—let's call it $e = (v_r, v_s)$—there are three possibilities:

- *(i) the edge e is already present in E_j.*
- *(ii) The edge $e = (v_r, v_s)$ is not in E_j, but the degrees of the vertices v_r and v_s are low in the sense that*

$$\deg_{G_j}(v_r) + \deg_{G_j}(v_s) < n,$$

where the subscript G_j is meant to show that the degree is being calculated in the graph G_j, whose vertex set is V and whose edge set is E_j. In this case we do not *include e in E_{j+1}.*

- *(iii) the edge $e = (v_r, v_s)$ is not in E_j, but the degrees of the vertices v_r and v_s are high in the sense that*

$$\deg_{G_j}(v_r) + \deg_{G_j}(v_s) \geq n. \tag{11.4}$$

Such an edge should be part of the closure, so we set

$$E_{j+1} \leftarrow E_j \cup \{e\}.$$

and then jump straight to step 3 below.

(3) Decide whether to stop: ask whether we added an edge during step 2.

- *If not, then stop: the closure $[G]$ has vertex set V and edge set E_j.*
- *Otherwise set $j \leftarrow j + 1$ and go back to step (2) to try to add another edge.*

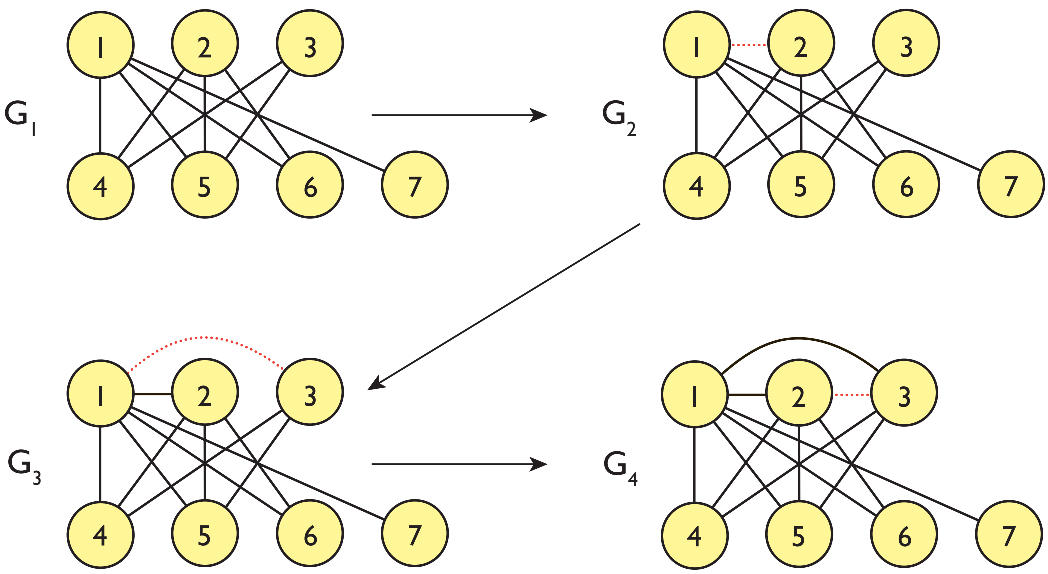

Figure 11.1: *The results of applying Algorithm 11.6 to the seven-vertex graph G_1. Each round of the construction (each pass through step 2 of the algorithm) adds a single new edge—shown with red, dotted curves—to the graph.*

11.2.2 An example

Figure 11.1 shows the result of applying Algorithm 11.6 to a graph with 7 vertices. The details of the process are discussed below.

Making G_2 from G_1
When constructing E_2 from E_1, notice that the vertex with highest degree, v_1, has $\deg_{G_1}(v_1) = 4$ and all the other vertices have lower degree. Thus, in step 2 of the algorithm we need only think about edges connecting v_1 to vertices of degree three. There are three such vertices—v_2, v_4 and v_5—but two of them are already adjacent to v_1 in G_1, so the only new edge we need to add at this stage is $e = (v_1, v_2)$.

Making G_3 from G_2
Now v_1 has $\deg_{G_2}(v_1) = 5$, so the closure condition (11.4) says that we should connect v_1 to any vertex whose degree is two or more, which requires us to add the edge (v_1, v_3).

Making G_4 from G_3
Now v_1 has degree $\deg_{G_3}(v_1) = 6$, so it is already connected to every other vertex in the graph and cannot receive any new edges. Vertex v_2 has $\deg_{G_3}(v_2) = 4$ and so should be connected to any vertex v_j with $\deg_{G_3}(v_j) \geq 3$. This means we need to add the edge $e = (v_2, v_3)$.

Conclusion: $[G] = G_4$
Careful study of G_4 shows that the rule (11.4) will not add any further edges, so the closure of the original graph is G_4.

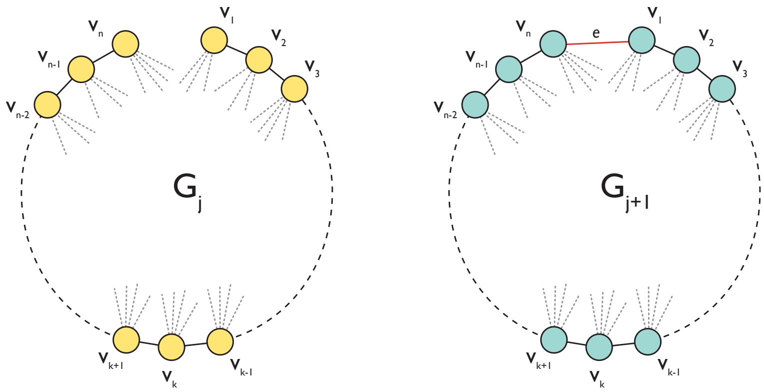

Figure 11.2: *The setup for the proof of the Bondy-Chvátal Theorem: adding the edge* $e = (v_1, v_n)$ *to* G_j *creates the Hamiltonian cycle* $(v_1, \ldots, v_n, v_1)$ *that's found in* G_{j+1}. *The dashed lines spraying off into the middles of the diagrams are meant to indicate that the vertices may have other edges besides those shown in black.*

11.3 The Bondy-Chvátal Theorem

The point of defining the closure is that it enables us to state the following lovely result:

Theorem 11.7 (Bondy and Chvátal, 1976)**.** *A graph G is Hamiltonian if and only if its closure $[G]$ is Hamiltonian.*

Before we prove this, notice that Dirac's and Ore's theorems are easy corollaries, for when $\deg(v) \geq n/2$ for all vertices (Dirac's condition) or when $\deg(u) + deg(v) \geq n$ for all non-adjacent pairs (Ore's condition), we have $[G] = K_n$, and, as we've seen, K_n is trivially Hamiltonian.

Proof. As the theorem is an if-and-only-if statement, we need to establish two things: (1) if G is Hamiltonian then $[G]$ is and (2) if $[G]$ is Hamiltonian then G is too. The first of these is easy in that the closure construction only *adds* edges to the graph, so in the sequence of edge sets (11.3) G has edge set E_1 and $[G]$ has edge set E_K with $K \geq 1$ and $E_K \supseteq E_1$. This means that any edges appearing in a Hamiltonian tour in G are automatically present in $[G]$ too, so if G is Hamiltonian, $[G]$ is also.

The second implication is harder and depends on an ingenious proof by contradiction. First notice that—by an argument similar to the one above—if some graph $G_{j_\star}$ in the sequence (11.2) is Hamiltonian, then so are all the other G_j with $j \geq j_\star$. This means that if the sequence is to begin with a non-Hamiltonian graph $G = G_1$ and finish with a Hamiltonian one $G_K = [G]$ there must be a single point at which the nature of the graphs in the sequence changes. That is, there must be some $j \geq 1$ such that G_j isn't Hamiltonian, but G_{j+1} is, even though G_{j+1} differs from G_j by only a single edge. This situation is illustrated in Figure 11.2, where I have numbered the vertices $v_1 \ldots v_n$ according to their position in the Hamiltonian cycle in G_{j+1} and arranged things so that the single edge whose addition creates the cycle is $e = (v_1, v_n)$.

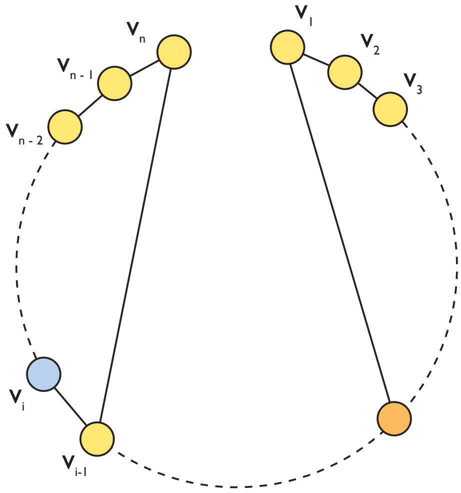

Figure 11.3: *Here the blue vertex, v_i, is in X because it is connected* indirectly *to v_n, through its predecessor v_{i-1}, while the orange vertex is in Y because it is connected* directly *to v_1.*

Let's now focus on G_j and define two interesting sets of vertices

$$X = \{v_i \mid (v_{i-1}, v_n) \in E_j \text{ and } 2 < i < n\}$$

and

$$Y = \{v_i \mid (v_1, v_i) \in E_j \text{ and } 2 < i < n\}.$$

The first set, X, consists of those vertices whose predecessor in the cycle has a direct connection to v_n, while the second set, Y, consists of vertices that have a direct connection to v_1: both sets are illustrated in Figure 11.3.

Notice that X and Y are defined to be subsets of $\{v_3, \ldots, v_{n-1}\}$, so they exclude v_1, v_2 and v_n. Thus X has $\deg_{G_j}(v_n) - 1$ members as it includes one element for each the neighbours of v_n except for v_{n-1}, while $|Y| = \deg_{G_j}(v_1) - 1$ as Y includes all neighbours of v_1 other than v_2. So then

$$\begin{aligned} |X| + |Y| &= \left(\deg_{G_j}(v_n) - 1\right) + \left(\deg_{G_j}(v_1) - 1\right) \\ &= \deg_{G_j}(v_n) + \deg_{G_j}(v_1) - 2 \\ &\geq n - 2 \end{aligned}$$

where the inequality follows because we know

$$\deg_{G_j}(v_n) + \deg_{G_j}(v_1) \geq n$$

as the closure construction is going to add the edge $e = (v_1, v_n)$ when passing from G_j to G_{j+1}. But then, both X and Y are drawn from the set of vertices $\{v_i \mid 2 < i < n\}$ which has only $n - 3$ members and so, by the pigeonhole principle, there must be some vertex v_k that is a member of both X and Y.

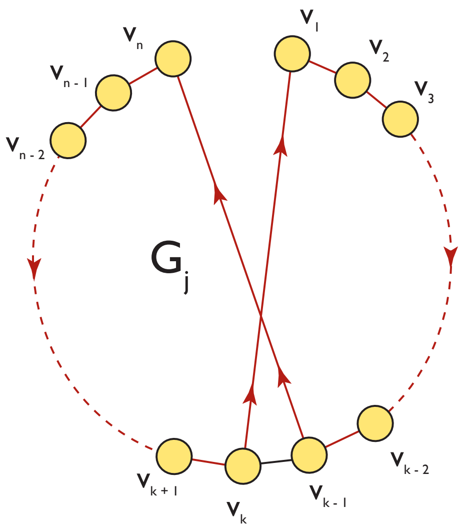

Figure 11.4: *The vertex v_k is a member of $X \cap Y$, which implies that there is, as shown above, a Hamiltonian cycle in G_j.*

The existence of such a v_k implies the presence of a Hamiltonian tour in G_j. As is illustrated in Figure 11.4, this tour:

- runs from v_1 to v_{k-1}, in the same order as the tour found in G_{j+1},
- then jumps from v_{k-1} to v_n: this is possible becase $v_k \in X$.
- The tour then continues, passing from v_n to v_{n-1} and on to v_k, visiting vertices in the opposite order from the tour in G_{j+1}
- and concludes with a jump from v_k to v_1, which is possible because $v_k \in Y$.

The existence of this tour contradicts our initial assumption that G_j is not Hamiltonian, but G_{j+1} is. This means no such G_j can exist: the sequence of graphs in the closure construction can never switch from non-Hamiltonian to Hamiltonian and so if $[G]$ is Hamiltonian, then G must be too. □

11.4 Afterword

Students sometimes have trouble remembering the difference between Eulerian and Hamiltonian graphs and I'm not unsympathetic: after all, both are named after very famous, long-dead European mathematicians. One way out of this difficulty is to learn more about the two men. Leonhard Euler, who was born in Switzerland, lived longer ago (1707–1783) and was tremendously prolific, writing many hundreds of papers that made fundamental contributions to essentially all of 18th century mathematics. He also lived in a very alien scientific world in that he relied on royal patronage, first from the Russian emperor Peter the Great and then, later, from Frederick the Great of Prussia and then finally, toward the end of his life, from Catherine the Great of Russia.

By contrast William Rowan Hamilton, who was Irish, lived much more recently (1805–1865). He also made fundamental contributions across the whole of mathematics—the distinction between pure and applied maths didn't really exist then—but he inhabited a much more recognisable scientific community, first working as a Professor of Astronomy at Trinity College in Dublin and then, for the rest of his career, as the directory of Dunsink Observatory, just outside the city.

Alternatively, one can remember the distinction between Eulerian and Hamiltonian tours by noting that everything about Eulerian graphs starts with 'E': *Eulerian* tours go through every *edge* in a graph and are *easy* to find. On the other hand, *Hamiltonian* tours go through every vertex and are *hard* to find.