

The role of migration traps in the formation of binary black holes in AGN disks

Maria Paola Vaccaro^{1,*}, Yannick Seif¹, and Michela Mapelli^{1,2,3,4}

¹ Universität Heidelberg, Zentrum für Astronomie (ZAH), Institut für Theoretische Astrophysik, Albert-Ueberle-Str. 2, 69120 Heidelberg, Germany

² Universität Heidelberg, Interdisziplinäres Zentrum für Wissenschaftliches Rechnen, Heidelberg, Germany

³ Physics and Astronomy Department Galileo Galilei, University of Padova, Vicolo dell'Osservatorio 3, I-35122 Padova, Italy

⁴ INFN, Sezione di Padova, Via Marzolo 8, I-35131 Padova, Italy

Received 5 August 2025 / Accepted 26 February 2026

ABSTRACT

Binary black holes (BBHs) forming in the accretion disks of active galactic nuclei (AGNs) represent a promising channel for gravitational-wave production. BBHs are often assumed to form at migration traps, i.e., radial locations where the Type I migration of embedded stellar-mass black holes (BHs) transitions from outward to inward. For this study, we tested this assumption by explicitly simulating the radial migration of BH pairs in AGN disks under different torque prescriptions, including thermal effects and the switch to Type II migration. We mapped where and when binaries form as a function of supermassive BH (SMBH) mass, disk viscosity, and migrating BH mass. We find that, for SMBH masses below $M_{\bullet}^{\text{thr}} \sim 10^8 M_{\odot}$, the majority of pair-up events occur near migration traps ($\geq 80\%$). In contrast, for higher SMBH masses, differential migration dominates and off-trap pair-ups can prevail. Certain disk configurations (e.g., $\alpha = 0.01$, $M_{\bullet} < 10^6 M_{\odot}$) present a significant overdensity of pair-ups even in the absence of traps due to traffic-jam accumulations where the gamma profile changes slope sharply. We also investigated hierarchical BBH formation, showing that higher-generation pair-ups cluster more tightly around trap or traffic-jam radii. Our results provide realistic prescriptions for BBH pair-up locations and timescales, highlighting the limitations of assuming fixed BBH formation sites.

Key words. black hole physics – gravitational waves – stars: black holes – stars: kinematics and dynamics – galaxies: active

1. Introduction

The direct detection of gravitational waves (GWs) by the LIGO and Virgo collaborations has revolutionized our ability to study compact object populations in the Universe, most notably binary black holes (BBHs, Abbott et al. 2023a,b). As the catalog of GW events continues to grow, a pressing challenge is to identify and characterize the astrophysical environments in which such binaries form and evolve.

One particularly promising environment for BBH formation is the accretion disk of an active galactic nucleus (AGN), where a gas-rich and high-density environment provides ideal conditions for dynamical interactions among black holes (BHs, e.g., McKernan et al. 2012, 2025; Bartos et al. 2017; Stone et al. 2017; Yang et al. 2019a; Secunda et al. 2020; Tagawa et al. 2020a,b, 2023; Samsing et al. 2022; Vaccaro et al. 2024; Cook et al. 2025; Delfavero et al. 2025; Dittmann et al. 2025; Xue et al. 2025). Indeed, stellar-mass BHs can become embedded in the gaseous disk surrounding a supermassive BH (SMBH, Wang et al. 2024; Rowan et al. 2025a; Whitehead et al. 2025) and thus interact gravitationally with the disk, experiencing migration torques analogous to those governing planetary migration in protoplanetary disks (Paardekooper et al. 2010; McKernan et al. 2012; Bellovary et al. 2016; Masset 2017; Grishin et al. 2024).

As BHs migrate radially within the disk, the probability of close encounters and subsequent pair-up into binaries is enhanced (DeLaurentiis et al. 2023; Li et al. 2023; Rozner et al.

2023; Rowan et al. 2023, 2024; Whitehead et al. 2024b,a). Furthermore, the dissipation of orbital energy through gas drag and GW emission can efficiently harden such binaries, leading to mergers within the AGN's lifetime (Peters 1964; Tiede et al. 2020; Ishibashi & Gröbner 2020, 2024; Rozner & Perets 2024; Rowan et al. 2025b).

BBH pair-up in AGN disks can occur either within migration traps – special radial locations where the total torque acting on a BH vanishes, halting its migration and promoting accumulation of several BHs – or throughout other regions of the disk, collectively referred to as “the bulk” (Wang et al. 2021). When two BHs migrate at different rates through the AGN disk, a process we refer to as differential migration, their radial paths may intersect. This increases the likelihood of close encounters that, under the influence of gas drag and GW emission, result in BBH formation. McKernan et al. (2020) find that more than 50% of BBH mergers originate in the bulk, but hierarchical mergers are efficient primarily within migration traps. Furthermore, Tagawa et al. (2020b) demonstrate that BBHs forming in the bulk through gas dynamical capture tend to migrate toward the migration trap as they harden.

Existing semi-analytical approaches focusing on the hierarchical merger scenario, including the FASTCLUSTER population synthesis tool developed by the authors (Mapelli et al. 2021, 2022; Vaccaro et al. 2024; Tornamenti et al. 2024), assign pair-ups to precomputed trap locations. While computationally convenient, this approach limits the physical realism of model outputs, especially for predicting the spatial distributions and formation times of merging binaries.

* Corresponding author: mariapaolavaccaro@gmail.com

For this study, instead, we explicitly integrated BH migration trajectories within AGN disks. We used the [Sirko & Goodman \(2003\)](#) model to describe the disk structure and employed the publicly available pAGN module ([Gangardt et al. 2024](#)) to compute self-consistent radial profiles of gas surface density, temperature, and pressure. We evaluated migration torques using both a classical Type I prescription ([Paardekooper et al. 2010](#); [Bellovary et al. 2016](#)) and an updated formulation including thermal effects ([Masset 2017](#); [Grishin et al. 2024](#)). We accounted for transitions to Type II migration under gap-opening conditions ([Lin & Papaloizou 1993](#); [Bryden et al. 1999](#); [McKernan et al. 2014](#); [Vaccaro et al. 2024](#); [Gilbaum et al. 2025](#)).

We simulated millions of independent BH pairs, drawn from a mass distribution generated with population-synthesis calculations ([Iorio et al. 2023](#)), and tracked their one-dimensional migration under gravitational torques. We accounted for differential migration rates between BHs of varying masses, and identified binary formation using a combination of Hill-stability and orbit-crossing criteria. We further incorporated hierarchical evolution by seeding higher-generation (Ng) BHs based on the merger outcomes of first-generation ($1g$) binaries, including recoil kicks ([Maggiore 2018](#)).

Our main objective is to quantify the spatial distribution of BBH pair-ups across a broad parameter space. This includes variations in SMBH mass, disk viscosity, torque prescription, and BH mass. Our results provide realistic prescriptions for pair-up radii and timescales, which can be incorporated into semi-analytic population synthesis codes to improve merger rate predictions and GW data interpretation.

2. Methods

2.1. Disk structure

We modeled the physical properties of the accretion disk using the widely accredited steady-state analytic [Sirko & Goodman \(2003\)](#) disk model (hereafter [SG](#)). We solved the corresponding one-dimensional equations self-consistently via the pAGN module for the Python programming language ([Gangardt et al. 2024](#)) for given input parameters such as the mass of the central SMBH, M_* , and the viscosity coefficient, α . We fixed the mass accretion rate to a tenth of the Eddington rate, $\dot{M}_* = 0.1 \dot{M}_{*,\text{Edd}}$. We obtained radial profiles for the gas surface density, Σ_{gas} , the disk aspect ratio, $h = H/R$ (defined as the ratio between the height of the disk, $H(R)$, and its radius, R), the temperature, T , and the sound speed, c_s .

Self-consistent solutions are crucial, as approximate treatments can lead to large deviations in these profiles, which in turn dramatically affect predictions for BBH formation in AGN disks, as discussed in [Appendix A](#).

2.2. Migration torques

We explicitly computed the torques felt by an embedded object of mass m at a radial distance R from a SMBH of mass M_* . We included the effects of Type I migration caused by the gravitational perturbation of embedded BHs in the disk ([Paardekooper et al. 2010](#); [Lyra et al. 2010](#); [Jiménez & Masset 2017](#)), accounting also for thermal migration, caused by the thermal response of the disk to the cold, overdense accretion disks surrounding embedded BHs ([Masset 2017](#)). The overall torque is

$$\Gamma = \Gamma_{\text{type I}} + \Gamma_{\text{th}}. \quad (1)$$

We investigated the difference between a classical Type I migration prescription, inspired by the work by [Bellovary et al. \(2016, hereafter B16\)](#), and one that also accounts for vertical and thermal effects, similar to [Grishin et al. \(2024, hereafter G24\)](#). The classical Type I migration torques were computed according to [B16](#) as

$$\Gamma_{\text{type I}} = \frac{\Gamma_{\text{ad}}\Theta^2 + \Gamma_{\text{iso}}}{(\Theta + 1)^2}, \quad (2)$$

where the parameter Θ was computed according to [Lyra et al. \(2010\)](#) as¹

$$\Theta = \frac{c_V \Sigma_g \Omega \tau_{\text{eff}}}{12\pi\sigma T^3}, \quad (3)$$

and the isothermal Γ_{iso} and adiabatic Γ_{ad} torques are functions of the adiabatic index γ (assumed to be equal to 5/3 throughout this work) and the gradients of local surface density Σ_g , and temperature T

$$a = -\frac{\partial \log \Sigma_g}{\partial \log R}, \quad b = -\frac{\partial \log T}{\partial \log R}, \quad c = b - (\gamma - 1)a, \quad (4)$$

as

$$\Gamma_{\text{iso}}/\Gamma_0 = -0.85 - a - 0.9b \quad (5)$$

$$\gamma\Gamma_{\text{ad}}/\Gamma_0 = -0.85 - a - 1.7b + 7.9c/\gamma. \quad (6)$$

[Jiménez & Masset \(2017\)](#) incorporated three-dimensional effects in their Type I migration torque,

$$\Gamma_{\text{type I}} = \Gamma_L + \left(\frac{0.46 - 0.96a + 1.8b}{\gamma} \right) \Gamma_0, \quad (7)$$

and also accounted for the Lindblad torque Γ_L as

$$\Gamma_L/\Gamma_0 = (-2.34 + 0.1a - 1.5b) f(\chi/\chi_c), \quad (8)$$

where

$$f(x) = \frac{(x/2)^{1/2} + 1/\gamma}{(x/2)^{1/2} + 1}, \quad \chi = \frac{16\gamma(\gamma - 1)\sigma_{\text{SB}}T^4}{3\kappa\rho_0^2H^2\Omega^2}, \quad \chi_c = H^2\Omega. \quad (9)$$

We took the electron-scattering opacity as $\kappa = 0.34 \text{ cm}^2/\text{g}$ as in [Gangardt et al. \(2024\)](#).

The thermal torques were computed as ([Masset 2017](#); [Velasco Romero & Masset 2020](#))

$$\Gamma_{\text{th}} = 1.61 \frac{\gamma - 1}{\gamma} \frac{x_c}{\lambda} \left(\frac{L}{L_c} \frac{4m_c}{m + 4m_c} - \frac{2m_c}{m + 2m_c} \right) \frac{\Gamma_0}{h} \left(1 - e^{-\lambda\tau_V/h} \right), \quad (10)$$

where $m_c = c_s\chi/G$ is a critical mass, τ_V is the vertical optical depth, x_c is the corotational radius of the embedded body and λ is the typical size of overdense regions surrounding embedded BHs, evaluated as in [G24](#):

$$x_c = -\frac{H^2}{3\gamma R} \frac{\partial \log P}{\partial \log R}, \quad \lambda = \sqrt{\frac{2\chi}{3\gamma\Omega}}. \quad (11)$$

We assumed that the embedded BH emits at Eddington luminosity, $L = L_{\text{Edd}} = 4\pi Gmc/\kappa$, and took $L_c = 4\pi Gm\rho\chi/\gamma$ as the critical luminosity ([Masset 2017](#)).

¹ Here c_V is the thermodynamic constant with constant volume, Ω is the Keplerian angular velocity, τ_{eff} is the effective optical depth, and σ is the Stefan-Boltzmann constant.

All migration torques are normalized by

$$\Gamma_0 = \left(\frac{m}{M_\bullet h} \right)^2 \Sigma_g R^4 \Omega^2. \quad (12)$$

We identified migration traps as locations in the disk where the total Type I migration torque Γ crosses zero, changing sign from positive to negative. These correspond to regions where BHs tend to accumulate. In contrast, locations where Γ changes sign from negative to positive are referred to as anti-traps, as they tend to be depleted of BHs.

We also accounted for the possibility of switching to Type II migration. Namely, we enforced a joint condition on the gas viscosity, migrator mass and local disk thinness, (Crida et al. 2006)

$$K = \frac{m^2}{\alpha h^5 M_\bullet^2} > 11. \quad (13)$$

If this condition is met, the migrator opens a gap in the disk and can no longer be subject to the migration torque Γ . Instead, conservatively assuming that no gas can cross the gap, its migration follows the viscous evolution of the disk's gas, moving inwards on a timescale (McKernan et al. 2012)

$$t_{\text{type II}} = t_{\text{visc}} = (\alpha h^2 \Omega)^{-1}. \quad (14)$$

Our prescription for the disk profile does not allow us to self-consistently open a gap (i.e., set $\Sigma_{\text{gas}}(R) \approx 0$ locally), hence in case of Type II migration we artificially constructed the corresponding torque as the rate of change in angular momentum, $L = mR^2\Omega$,

$$\Gamma_{\text{type II}} = \left. \frac{dL}{dt} \right|_{\text{type II}} \simeq \frac{L}{t_{\text{type II}}} = \alpha m H^2 \Omega^2. \quad (15)$$

We discuss the validity and impact of this assumption in Appendix B.

2.3. Integration scheme and pair-up criterion

We began each simulation by selecting two stellar-mass BHs that were assumed to already be embedded in the AGN disk. For stellar-origin BHs, henceforth referred to as first generation (1g) BHs, masses m are sampled independently from a distribution obtained via population-synthesis modeling SEVN for solar metallicity (Iorio et al. 2023). The initial radial location R_{mit} , is drawn from a uniform distribution in logarithmic space. Specifically, the logarithm of the initial radius is uniformly sampled between the logarithm of the disk's inner radius, assumed to be at the innermost stable circular orbit, $R_{\text{min}} = 6R_g$ in this context, and an outer radius $R_{\text{max}} = 0.1 \text{ pc } (M_\bullet / 10^6 M_\odot)^{1/2}$, beyond which the disk's self-gravity becomes important (Goodman 2003; Yang et al. 2019b). This choice reflects the expectation that BHs can form or become embedded at various locations throughout the AGN disk, and provides a neutral baseline that does not bias our results toward any particular radial scale. However, the stellar-origin BHs we expect to find in an AGN disk come from the surrounding NSC, and the probability that they are captured by the disk depends on the properties of both the AGN disk and the NSC. As a result, embedded BHs are expected to accumulate preferentially in particular radial ranges, enhancing the probability of binary formation at those locations relative to our log-uniform assumption.

The subsequent orbital evolution is treated in one dimension, considering only the radial location as a function of time. The

equation of motion is derived from angular momentum conservation in a Keplerian potential and takes the form:

$$\dot{R} = \frac{2\Gamma(R)}{m} \sqrt{\frac{R}{GM_\bullet}}, \quad (16)$$

where $\Gamma(R)$ is obtained as in Eq. (1) or (15) depending on whether the conditions for Type II migration (Eq. 13) are met. The torque profile is precomputed from the AGN disk model and interpolated at each timestep to reflect the actual mass and location of the migrating body.

We solved this system of coupled differential equations using the second-order Runge–Kutta integration method with embedded third-order error estimation and an adaptive timestep. The integration spans the full duration of the AGN disk's activity, which we fix at 2 Myr, and is initialized with a timestep of 10^3 yr. To capture statistical trends in binary formation, this procedure was repeated for a large number of random initial conditions. We simulated 10^6 independent realizations, each with different initial masses and orbital radii, resulting in a comprehensive catalog of pair-up events across a wide range of physical parameters.

A central goal of the simulation is to determine whether and when a given pair of migrating single BHs becomes a bound binary. To that end, we randomly select two BHs in our catalog and we label them as primary (mass m_1 , radial location R_1) and secondary (m_2, R_2), with the convention that $m_1 \geq m_2$. We applied two complementary criteria for binary formation in post-processing, after completing the orbital integration of each single BH.

The first and most stringent condition is based on Hill stability (Hill 1878). We computed the instantaneous separation between the two BHs, $d = |R_1 - R_2|$, at each integration step and compare it to their Hill radii. A binary is considered to have formed if the separation falls below the larger of the two Hill radii, hence the two BHs are close enough for their mutual gravitational attraction to overcome tidal forces from the central SMBH:

$$d \leq \max_{i=1,2} R_H(m_i, R_i), \quad R_H(m, R) = R \left(\frac{m}{3(M_\bullet + m)} \right)^{1/3}. \quad (17)$$

To prevent the integrator from skipping over close encounters due to large timesteps or rapid migration, we also monitored the sign of the radial separation $E_{\text{cross}} = R_1 - R_2$. A sign change in E_{cross} indicates that the BHs have crossed each other's orbits and may have come into sufficiently close proximity to allow for capture. If either the explicit Hill condition or this indirect orbit-crossing condition is met, a BBH is considered to have formed.

In either case, we recorded the BBH formation time and its pair-up radius, defined as the mass-weighted average position of the two BHs at the moment of binary formation:

$$R_{\text{pair}} = \frac{m_1 R_1 + m_2 R_2}{m_1 + m_2}. \quad (18)$$

This dual-criterion strategy balances accuracy and robustness, ensuring that bound systems are identified even when the BHs migrate rapidly or interact near the resolution limit of the integrator. However, it neglects the exact binary binding mechanism under the influence of gas torques and disk turbulence, investigated in hydrodynamical simulations (e.g. Whitehead et al. 2024b). The resulting catalog of binary formation events provides the foundation for our statistical analysis of BBH assembly in AGN disks.

2.4. *N*th-generation BHs

For *N*th-generation (*Ng*) BHs, which are the remnants of earlier mergers in the AGN disk, we estimated their orbital starting conditions using the outputs of *1g* simulations. Hence, in the *Ng* simulations, the initial position R_{init} of the primary BH is no longer drawn randomly across the disk. Instead, it is sampled directly from the set of R_{pair} values obtained in the *1g* simulation, accounting for the natal velocity kick v_{kick} that is imparted during the merger event. The recoil velocity v_{kick} is drawn from a precomputed FASTCLUSTER catalog (Vaccaro et al. 2024) that already incorporates the dependence on progenitor mass ratio and spin configuration as in Maggiore (2018). The kick direction is randomized assuming isotropic ejection, and contributions from enclosed gas and NSC members are neglected given their small mass relative to the SMBH. If the kick magnitude exceeds the local escape velocity, the remnant is assumed to leave the system. Otherwise, the object is recaptured and is then able to migrate in the disk following the same torque prescriptions as *1g* BHs. The shifted radius $R_{\text{init}}^{\text{new}}$ is computed by means of simple orbital transfer calculations (Hohmann 1960) as

$$R_{\text{init}}^{\text{new}} = \frac{GM_{\bullet}}{v_{\text{tot}}^2} \left[\frac{R_{\text{init}} v_{\text{tot}}^2 / GM_{\bullet}}{2 - R_{\text{init}} v_{\text{tot}}^2 / GM_{\bullet}} \right], \quad (19)$$

where v_{tot} is the vector sum of the circular orbital velocity $v_{\text{Kepl}} = \sqrt{GM_{\bullet}/R}$ and the random kick v_{kick} .

In doing so, we implicitly assumed that *1g* binaries will eventually merge at a radial location not significantly different from where they originally paired up.

We only allowed *Ng* BHs to pair up with first-generation companions (forming *Ng*–*1g* BBHs). This choice is motivated by the low probability of drawing a secondary of generation M_2 given the steeply decreasing $p(M) \propto 2^{-(M-1)}$ (Zevin & Holz 2022). Accordingly, the secondary BH in each pair has mass and initial radial position sampled in the same way as in the *1g* simulations.

This methodology ensures that *Ng* BHs are seeded in dynamically consistent locations, and allows us to model the spatial memory and redistribution of hierarchical mergers across the AGN disk. Other than the initial radius, the rest of the computation proceeds in an identical way to what previously described for *1g* BHs.

2.5. Description of runs

We performed a suite of Monte Carlo simulations, running $N = 10^5$ realizations for each combination of model parameters. We explore how BBH formation in AGN disks depends on the mass of the central SMBH, the disk’s viscosity, and the underlying migration prescription.

We vary the SMBH mass over the range $\log M_{\bullet}/M_{\odot} = 5.0$ – 9.0 , encompassing typical values expected across cosmic time (e.g. Greene & Ho 2007; Trinca et al. 2022).

The disk viscosity for a *SG* disk, and more generally for so-called α -disk models, is parametrized by coefficient $\alpha \in [0, 1]$. We adopt $\alpha = 0.1$ and $\alpha = 0.4$ to bracket the range suggested by observations of fully ionized accretion disks, which typically yield $\alpha \sim 0.1$ – 0.4 (King et al. 2005; Martin et al. 2019). We also ran some models with $\alpha = 0.01$ in order to align with conventional assumptions in theoretical models that often presume lower turbulence efficiency (e.g. standard disk models such as *SG*) and magneto-hydrodynamical simulations (e.g. Hawley et al. 2011). We are limited by the assumption of a

spatially constant α across the disk radius, consistent with the approximation in the *SG* model, despite more recent simulations indicating potential radial variation in viscosity (e.g. Penna et al. 2012).

To investigate the impact of different assumptions on the migration mechanism, we distinguished between a *B16* prescription, where we set $\Gamma_{\text{th}} = 0$ in Eq. (1) and computed $\Gamma_{\text{type I}}$ as in Eq. (2) (as in Bellovary et al. 2016), and a *G24* prescription where we set Γ_{th} as in Eq. (10) and $\Gamma_{\text{type I}}$ as in Eq. (7) (as in Grishin et al. 2024).

3. Results

3.1. Migration torques and pair-up radii

We begin by illustrating the torque structure that governs BH migration in *SG* disks. Figure 1 shows representative profiles of the total torque acting on a BH as a function of radial distance.

The torque intensity and sign vary depending on the local properties of the disk such as temperature and density gradients as in Eqs. (2)–(9), which are themselves functions of α , R , and M_{\bullet} . The mass of the migrator m only affects the intensity of the torque via the normalization factor in Eq. (12). Hence, migration trap locations are disk-specific and independent on the characteristics of the migrating BHs. Multiple torque sign reversals from $\Gamma > 0$ to $\Gamma < 0$ are generally present, corresponding to migration traps, whose number and position are of course also sensitive to M_{\bullet} and α .

To connect this torque structure with the outcome of binary formation, we show in Figure 2 the full distribution of BBH pair-up locations across our catalog of simulations. Each panel displays a two-dimensional histogram in the radius–SMBH mass plane, indicating the logarithmic number of pair-ups in each bin. We also indicate the locations of migration traps, as extracted from the torque profiles. Most BHs that successfully form binaries do so while migrating under Type I migration, with fewer than 8% forming under Type II conditions.

We find that the radial distribution of pair-ups is broadly consistent with the location of migration traps, supporting the idea that traps are favorable sites for binary formation. However, the distribution shows a significant spread around traps as well as some clear offsets, with efficient pair-ups also occurring at locations not identified as traps, while some traps contain no BBH pair-ups at all. These departures can be attributed to differential migration: mismatches in migration speed may cause pair of BHs to pair up upstream of the trap radius.

In some cases, such as at $R \approx 10^3 R_S$ for $M_{\bullet} < 10^6 M_{\odot}$ in the *B16* prescription, there is a clear overdensity of pair-ups at locations which are not marked as traps, caused by sharp changes in the slope of Γ . We refer to these locations as traffic jams.

In systems with smaller SMBHs ($\log M_{\bullet}/M_{\odot} \lesssim 6$), pair-ups are possible at almost all radial locations in the disk, whereas at higher SMBH masses ($\log M_{\bullet}/M_{\odot} \gtrsim 7$), the distribution of pair-up radii becomes significantly narrower and more concentrated toward the inner disk regions ($10^2 \lesssim R/R_S \lesssim 10^4$). This narrowing reflects the increased gravitational influence of more massive SMBHs which dominates over local disk perturbations, leading to a more stable torque structure. In particular, the radial torque Γ at large radii is less intense and uniformly negative at high M_{\bullet} , reducing differential migration between BHs and suppressing their pairing in the outer disk region.

We observe relevant differences in the trap location based on the torque prescription, as shown in Figure 2. Using the full *G24* migration torque gives rise to outer migration traps which are

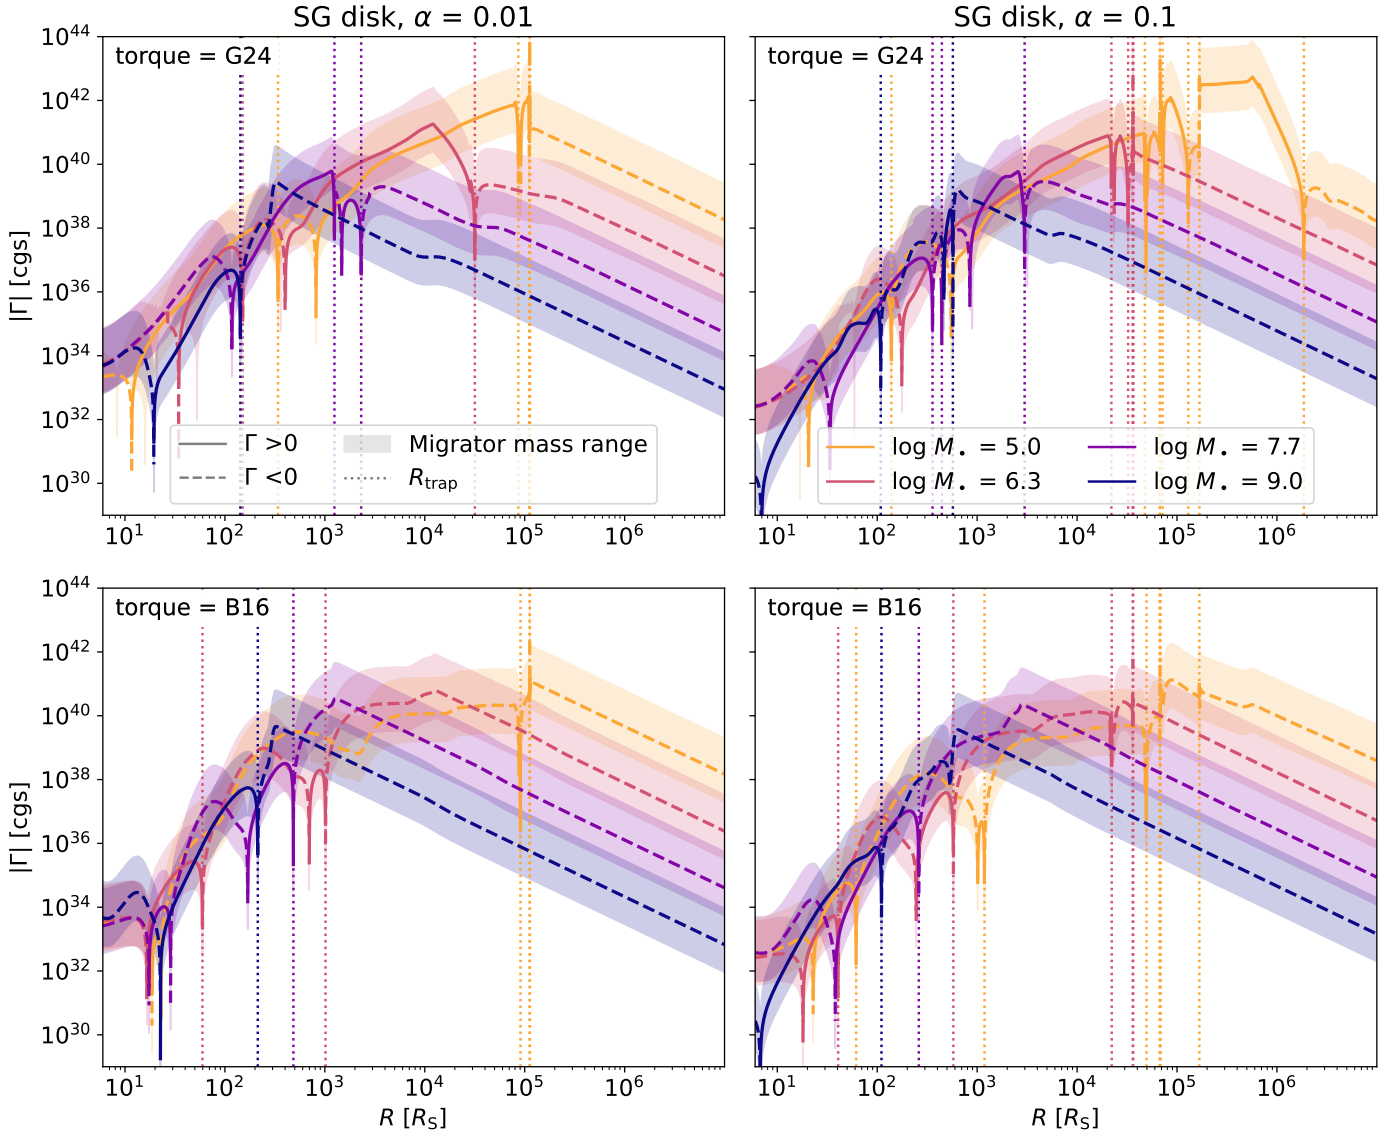


Fig. 1. Examples of torque profiles as a function of radial distance (in Schwarzschild radii) for BHs embedded in a SG disk. Each plot corresponds to a different combination of torque prescription (top: G24; bottom: B16), disk viscosity parameter α (left: $\alpha = 0.01$; right: $\alpha = 0.1$) and SMBH mass ($\log M_{\bullet}/M_{\odot} = 5.0, 6.3, 7.7, 9.0$, represented in different colors). Solid and dashed lines indicate positive and negative net torques, while shaded regions denote uncertainty in the torque determined by computing it over a range of migrator masses between $5 M_{\odot}$ and $50 M_{\odot}$. Vertical dotted lines mark the locations of migration traps.

absent in the classical B16 case for SMBH masses $\log M_{\bullet}/M_{\odot} \in [5.6, \sim 8.0]$ ($[6.7, 9.0]$) for $\alpha = 0.01$ (0.1).

3.2. The relevance of migration traps

To quantify the role of migration traps in BBH formation, we track the radial locations at which BHs pair up and compare them to the positions of predefined migration traps in the disk. For each simulation, we identify a set of trap locations $\{R_{\text{trap}}\}$ based on the local torque profile, as described in Section 2.2. A pair-up event is considered to occur “within a trap” if the radial distance between the pair-up location R_{pair} and any trap satisfies the following logarithmic tolerance condition:

$$|\log R_{\text{pair}} - \log R_{\text{trap}}| < \delta, \quad (20)$$

where δ is a fixed tolerance. In this work, we adopt $\delta = 0.05$.

Based on this definition, we compute the fraction f_{trap} as

$$f_{\text{trap}} = \frac{N_{\text{in trap}}}{N_{\text{pair}}}, \quad (21)$$

where $N_{\text{in trap}}$ is the number of pair-up events satisfying the condition in Eq. (20), and N_{pair} is the total number of pair-up events recorded in the simulation.

Figure 3 shows the resulting fraction f_{trap} , plotted as a function of the central SMBH mass. Overall, the majority of pair-ups occur within migration traps across most of the parameter space explored. For SMBH masses below a certain threshold M_{\bullet}^{thr} , f_{trap} has values of 0.7 to 0.9 for $1g$ BHs and it approaches unity for Ng BHs. As the SMBH mass increases, f_{trap} presents a pronounced transition from values close to unity to nearly zero, indicating that pair-up events are no longer connected with trap locations at high SMBH masses. This transition typically takes the form of




Fig. 2. Two-dimensional histograms showing the distribution of BBH pair-up locations in the radius–mass plane (R vs. M_{\bullet}), for different combinations of torque prescription (top: **G24**; bottom: **B16**) and disk viscosity α (left: 0.01; right: 0.1) in a **SG** disk. The color map indicates the number of pair-ups in each bin N_{pair} . Red crosses mark the locations of migration traps.

an abrupt drop from $f_{\text{trap}} \simeq 1$ to $f_{\text{trap}} \simeq 0$, with specific value of M_{\bullet}^{thr} depending on the disk and torque model adopted.

In the high-viscosity **B16** models ($\alpha = 0.1, 0.4$), the drop occurs early, at $M_{\bullet}^{\text{thr}} \simeq 10^{7.4} M_{\odot}$, whereas in the low-viscosity case ($\alpha = 0.01$), $M_{\bullet}^{\text{thr}} \simeq 10^{8.5} M_{\odot}$. For the **G24** model, the transition to low f_{trap} consistently occurs at large SMBH masses; specifically, at $M_{\bullet}^{\text{thr}} \simeq 10^{8.8} M_{\odot}$, $>10^9 M_{\odot}$, $\simeq 10^{8.5} M_{\odot}$ for the cases with $\alpha = 0.01, 0.1, 0.4$. Furthermore, f_{trap} curves present some outliers due to peculiar disk properties for specific combinations of M_{\bullet} and α (see Figure 2 and additional $|\Gamma|$ plots on Zenodo). For instance, in disks with $\log M_{\bullet}/M_{\odot} \in [5.0, 6.0]$ and $\alpha = 0.01$, f_{trap} increases gradually from 0.2 to ~ 0.8 for the **B16** torque due to traffic jams: a sharp transition in the slope of Γ gives rise to many pair-up events outside of traps due to differential migration. Specifically, as visible in Figure 1, this happens at $R \lesssim 10^2 R_S$ for $\log M_{\bullet}/M_{\odot} = 5.0\text{--}5.1$ and at $R \sim 10^3 R_S$ for $\log M_{\bullet}/M_{\odot} < 6.0$.

We show normalized histograms of binary pair-up radii R_{pair} , split by BH generation ($1g$ vs Ng), in Figure 4 for selected SMBH

masses and disk viscosity. Pair-up locations for Ng BHs are considerably more concentrated around peaks than for $1g$ BHs. This enhanced localization arises from the initial conditions: in our model, Ng remnants resume migration at a radius set by Eq. (19). Because recoil speeds v_{kick} are typically small relative to the local Keplerian velocity v_{Kepl} (Maggiore 2018; Vaccaro et al. 2024), the kick induces only a minor radial displacement, so Ng remnants remain close to the pair-up radius of their progenitor binary. Since $1g$ binaries preferentially form near peaks in the R_{pair} distribution, the ensuing Ng BHs tend to start their migration near those same radii. As a result, R_{pair} for Ng objects clusters tightly around peaks, which naturally explains the systematically higher values of f_{trap} seen for Ng in Figure 3 since most peaks are coincident with traps. In contrast, $1g$ BHs are initialized with randomly distributed radial positions within the disk, and their subsequent migration is governed entirely by the local torque profile. While traps still influence the dynamics, the broader range of initial conditions results in a more diffuse distribution of pair-up radii, and therefore a lower f_{trap} .




Fig. 3. Fraction of binary pair-ups occurring within migration traps as a function of the SMBH mass. The vertical axis shows the occurrence f_{trap} of pair-up events that take place in proximity to predefined migration trap locations, while the horizontal axis shows the logarithm of the central SMBH mass. Different marker fillings indicate the BH generation (filled: $1g$; hollow: Ng), while color encodes the viscosity parameter α (orange: 0.01; teal: 0.1; navy blue: 0.4). The two torque formalisms are distinguished by marker shape and line style (squares and solid lines: G24; circles and dashed lines: B16). Trap proximity is defined using a logarithmic tolerance on radial separation (see Section 3.2 for details).

For low viscosity ($\alpha = 0.01$) under the B16 prescription, we also find a pronounced pile-up of pair-ups at $R \simeq 10^3 R_S$ for $M_* \lesssim 10^6 M_\odot$ that does not coincide with migration trap locations. This is a direct consequence of differential migration: as shown in Figure 1, the $|\Gamma(R)|$ profiles exhibit an abrupt change in slope at these radii, producing a traffic-jam accumulation of pair-ups.

These findings emphasize the importance of migration traps as anchoring points for BBH formation in AGN disks. Their effectiveness is sensitive to both disk properties (through the viscosity parameter α) and central SMBH mass, and is particularly crucial for Ng pair-ups and their resulting hierarchical mergers.

3.3. Pair-up efficiency

We define the pair-up efficiency η as the fraction of BHs in a given simulation that successfully form a binary by the end of the AGN disk lifetime. That is,

$$\eta = \frac{N_{\text{pair}}}{N_{\text{tot}}}, \quad (22)$$

where N_{pair} is the number of BHs that participated in a pair-up event, and N_{tot} is the total number of simulated BHs. Figure 5 shows the behavior of η as a function of the central SMBH mass, for both $1g$ and Ng BHs. As expected, the efficiency tends to decrease with increasing SMBH mass, due to the suppression of differential migration in the outer regions of the disk discussed in Section 3.1. Nonetheless, Ng BHs consistently exhibit higher pair-up efficiency than their $1g$ analogs. This is a consequence of their more favorable initial conditions as previously observed for f_{trap} : they typically start closer to trap or traffic-jam locations, which increases their likelihood of encountering a suitable companion before the end of the disk lifetime.

We also observe a systematic, although modest, dependence of η on the value of α and on the assumed migration model. For low-viscosity disks ($\alpha = 0.01$), pair-up efficiency remains higher across a broader range of SMBH masses, whereas higher values of α (0.1, 0.4) slightly reduce the likelihood of BBH formation. Interestingly, these differences are more pronounced for

the Ng BHs than for $1g$, where the impact of α is comparatively marginal. This may be due to the inherently stochastic nature of pair-up dynamics in the outer disk regions at high viscosity ($\alpha \geq 0.1$), relatively low SMBH masses ($\log M_*/M_\odot \lesssim 6.5$), where multiple nearby trap locations (noticeable in Figures 1 and 2) can lead to complex migration patterns.

We also note marked jumps in η which correlate with the temporary disappearance of outer migration traps in the corresponding M_* range, which allows outer BHs to move toward intermediate-radius ($\sim 10^3 R_S$) locations, characterized by sudden changes in the slope of Γ which facilitate BBH captures by differential migration. This happens most notably at $\log M_*/M_\odot \sim 5.6-5.8$ for $\alpha = 0.01$, and at $\log M_*/M_\odot \gtrsim 6.7$ for $\alpha = 0.1$.

3.4. Impact of initial conditions

Figure 6 is an analysis of how the pair-up efficiency η and the trap fraction f_{trap} depend on the initial conditions of the migrating BH, specifically its initial radial position R_{init} and mass m .

The pair-up efficiency η shows a strong radial dependence, peaking at intermediate radii ($R_{\text{init}} \sim 10^2-10^5 R_S$). This trend reflects the underlying structure of the torque profile: migration is inefficient near the inner and outer edges of the disk, while the central regions allow for efficient dynamical pairing. In the innermost disk, pair-ups are suppressed by limited radial volume and weaker torques, whereas in the outer disk, BHs typically migrate too slowly to encounter a companion within the AGN lifetime. A secondary dependence on mass is also evident: heavier BHs ($m \gtrsim 20 M_\odot$) tend to achieve higher pair-up efficiencies due to their stronger interaction with the gas and faster migration rates.

The trap fraction f_{trap} follows a similar, but more pronounced, trend to the pair-up efficiency η in terms of BH mass m . In particular, stellar-mass BHs with $m \lesssim 20 M_\odot$ are generally unable to reach migration traps, with average f_{trap} remaining below 0.3, whereas BHs with $m \gtrsim 20 M_\odot$ consistently exhibit trap fractions exceeding 0.6 (0.9) in the B16 (G24), indicating that higher-mass objects are not only more efficient migrators but also more likely to pair up in the vicinity of traps.

However, the dependence of f_{trap} on the initial radius R_{init} appears more complex than that of the overall efficiency η . While the peak values of both quantities occur broadly in the same intermediate radial range ($\sim 10^2-10^5 R_S$), the trap fraction displays sharper and more localized peaks. This suggests that proximity to a trap at birth plays a significant role: BHs that happen to become embedded sufficiently close to a stable trap are more likely to remain in its vicinity and undergo in-trap pair-up. Consequently, f_{trap} is subject to a degree of stochasticity tied to the initial orbital configuration, which is less evident in the smoother behavior of η .

3.5. Migration timescales

In semi-analytical models of BH dynamics in AGN disks, the migration timescale is often estimated using the rate of angular momentum change via gaseous torques. A common approximation (e.g. McKernan et al. 2012; Vaccaro et al. 2024) is to compute the migration time as the ratio between the angular momentum L of the BH and the torque normalization Γ_0 :

$$t_{\text{migr},0} = \frac{L}{\Gamma_0} = \frac{M_*^2 h^2}{m \Sigma_g R^2 \Omega} \quad (23)$$




Fig. 4. Histograms of binary pair-up radii R_{pair} (in units of Schwarzschild radii) for first-generation ($1g$) and N th-generation (Ng) BHs, across a range of SMBH masses ($\log M_*/M_\odot = 6, 7, 8$ in the three columns, respectively) and viscosity parameters (upper row: $\alpha = 0.01$; lower row: $\alpha = 0.1$). Panels show the pair-up radii probability density functions (PDFs), normalized separately in each panel. The $1g$ distributions are shown as step histograms, while Ng distributions are overlaid either as filled or hatched histograms for visual distinction. Vertical dotted lines indicate the positions of migration traps identified from the torque profile. Color represents the migration model used (green: G24; purple: B16).




Fig. 5. Pair-up efficiency as a function of the central SMBH mass. The vertical axis shows the fraction of BHs that successfully form a binary during the simulation, while the horizontal axis shows the logarithm of the SMBH mass. As in Figure 3, different marker fillings indicate the BH generation (filled: $1g$; hollow: Ng), color encodes the viscosity parameter α (orange: 0.01; teal: 0.1; navy blue: 0.4), and the two torque prescriptions are distinguished by marker shape and line style (squares and solid lines: G24; circles and dashed lines: B16).

where $L = mR^2\Omega$ and Γ_0 is as in Eq. (12). This simplified expression is widely used in the literature as it captures the

characteristic timescale over which a BH would drift through the disk under idealized conditions. In Figure 7, the resulting timescales computed for a $10 M_\odot$ BH are shown as dashed lines for different SMBH masses, viscosity parameters, and torque prescriptions.

However, when using the full torque Γ (Eq. 1) the corresponding migration timescale, $t_{\text{migr}} = L/\Gamma$, deviates from the Γ_0 -based estimate, especially in regions near torque sign-reversal points, i.e., in the vicinity of migration traps or anti-traps, where $\Gamma = 0$. As a result, the solid lines in Figure 7 exhibit sharp features and even divergences, reflecting the breakdown of the timescale approximation in those zones.

To assess how well either of these approximations tracks the actual dynamical evolution of BHs in the disk, we compare them to the pair-up times, t_{pair} , obtained directly from our simulations. These pair-up times, shown via the 40% and 68% contour levels in Figure 7, represent the elapsed time between the moment at which the single primary BH becomes embedded in the disk and the subsequent formation of a BBH. Unlike t_{migr} , they are measured durations rather than derived timescales, and naturally include the effects of stochasticity, mass differences, and local disk structure. Several important features emerge from this comparison. There is substantial scatter in the pair-up times at fixed radius, due in part to the variation in the mass of the migrating BH. There is a sharp upper limit at $t_{\text{migr}} = 2 \text{ Myr}$, corresponding to the assumed AGN disk lifetime adopted in our

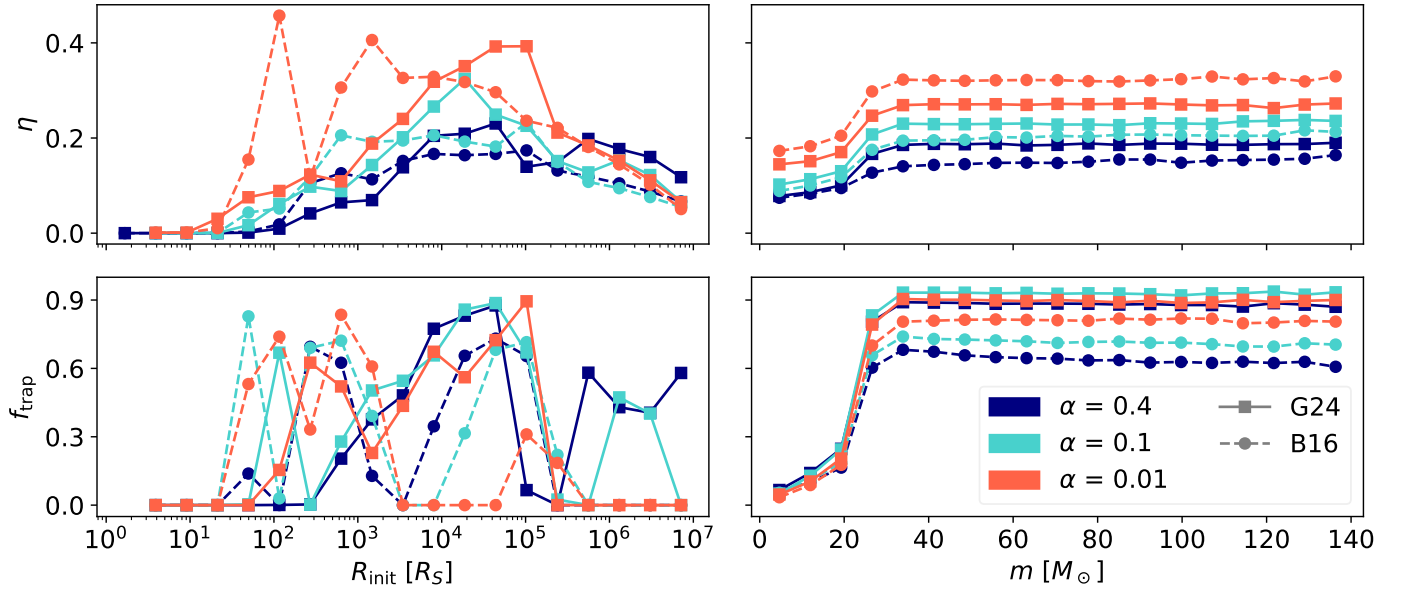


Fig. 6. Pair-up efficiency (top panels) and fraction of pair-ups occurring within migration traps (bottom panels) as functions of initial radius R_{init} (left) and migrator mass m (right). Color encodes the viscosity parameter α (orange: 0.01; teal: 0.1; navy blue: 0.4), and the two torque prescriptions are distinguished by marker shape and line style (squares and solid lines: G24; circles and dashed lines: B16).

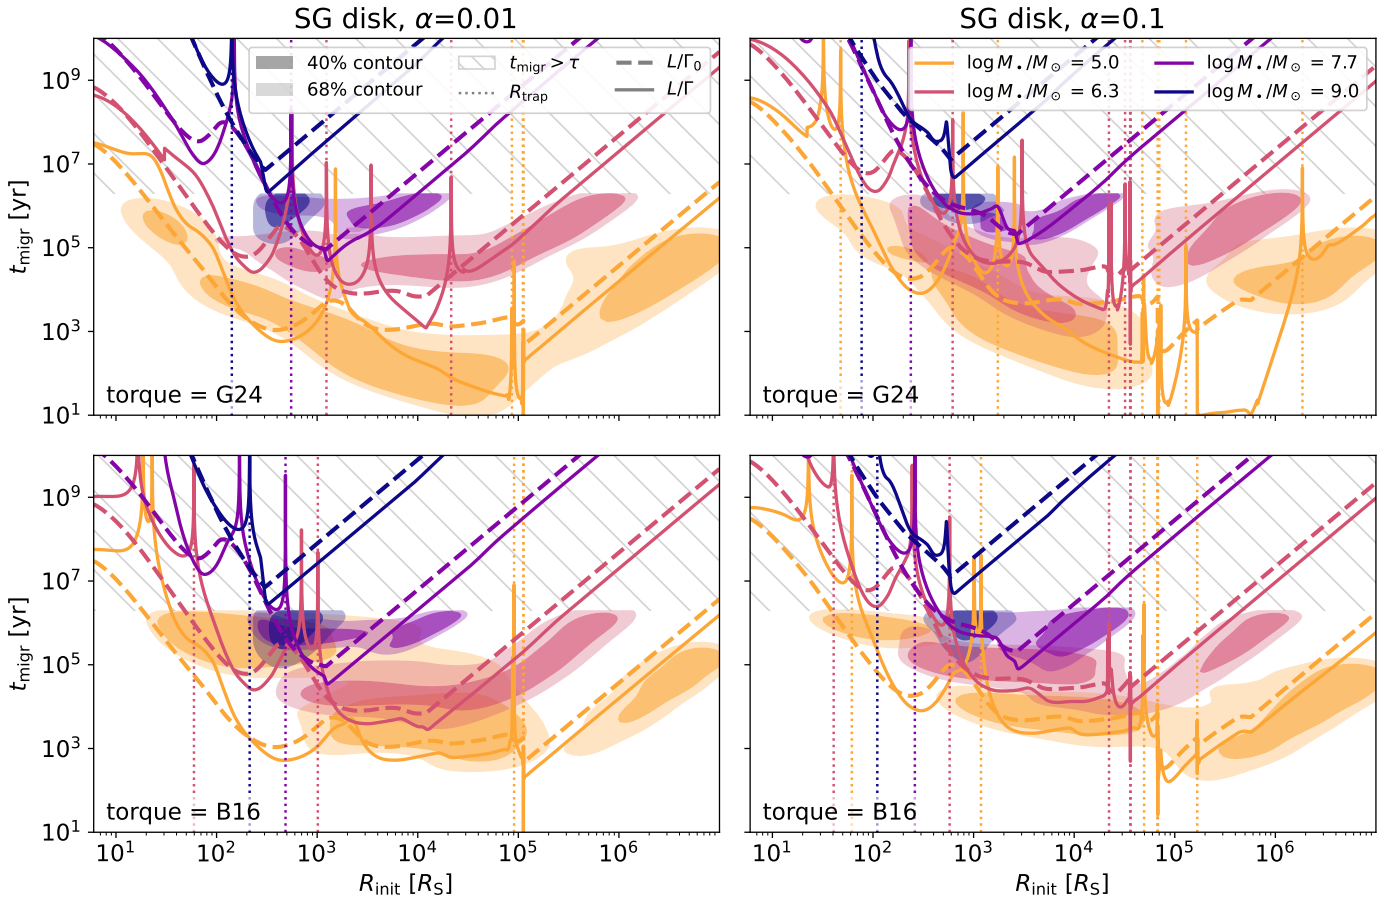


Fig. 7. Two-dimensional distributions of the migration time t_{migr} as a function of the initial radius R_{init} in SG AGN disks, for different combinations of the viscosity parameter α (left: 0.01; right: 0.1) and torque prescriptions (top: G24; bottom: B16). Each panel shows contours enclosing 40% and 68% of the migration time distribution, based on our simulations. Results are shown for four representative SMBH masses: $\log(M_*/M_\odot) = 5.0, 6.3, 7.7, 9.0$, indicated by color. Solid (dashed) lines represent the migration timescale computed using the torque Γ in Eq. (1) (Γ_0 in Eq. 12). Vertical dotted lines mark the locations of migration traps.

model. Moreover, some radial regions are systematically disfavored for binary formation, as also noticeable in Figure 6, due to either negative torques in the innermost regions of the disk or to reduced differential migration in the outer disk regions for large M_\bullet , as discussed in Section 3.1. Finally, pair-up times in our simulations are typically shorter than the estimated migration timescale in regions where the Γ profile has a large slope, and longer where the profile is flatter. This is a consequence of differential migration: pair-ups are set by relative drift rather than absolute torque. Where the migration speed profile is steep, nearby objects converge quickly and pair up fast. Vice versa, where the profile is flat, velocity differences are only mass-dependent and consequently pair-up times lengthen.

4. Caveats

Our current approach still presents several limitations.

First, we do not model gas accretion onto individual BHs. Accretion is expected to lead to spin alignment and mass growth, which in turn influences migration rates and pair-up probabilities. We also neglect any vertical stratification in the AGN disk and restrict all BHs to move strictly in the midplane, thereby omitting any effect related to vertical motion. In addition, the SG model is not fully realistic, as it assumes an unspecified heating source and does not account for gas consumption even in gravitationally unstable regions.

Our treatment of Type II migration is deliberately conservative, as gap-opening objects are assumed to effectively prevent any inflow across the gap. More realistically, the transition to Type II is not a hard switch: the gap would not open fully and gas inflow would not be completely suppressed (Malik et al. 2015). As a result, the Type II regime would more closely resemble a modified Type I migration with reduced local density (as shown in Kanagawa et al. 2018; Fung et al. 2014). We explore this alternative model for Type II migration in Appendix B.

Moreover, our model does not include resonant interactions, such as mean-motion resonances, which could alter or halt migration (Epstein-Martin et al. 2025; Moncrieff et al. 2026). We also do not consider dynamical formation via three-body scatterings, which is another viable mechanism of BBH formation in AGN environments (Tagawa et al. 2020b). Finally, while our pair-up criterion identifies candidate binaries based on proximity, it does not explicitly model binary binding or survivability under the influence of gas torques and disk turbulence.

We will cure these approximations in a forthcoming study.

5. Summary

Semi-analytical models such as FASTCLUSTER (Mapelli et al. 2021, 2022; Vaccaro et al. 2024; Tornamenti et al. 2024) and McFACTS (McKernan et al. 2025; Cook et al. 2025; Delfavero et al. 2025) have proven to be useful tools for modeling hierarchical mergers of BHs in AGN disks. Their computational efficiency allows for rapid exploration of the parameter space and facilitates comparisons with observational data. However, these models often rely on simplified assumptions regarding the locations at which BBHs form. A common assumption, made for example by Vaccaro et al. (2024), is that all pair-ups occur at so-called migration traps. While this assumption captures part of the physical picture, it overlooks the diversity of pair-up locations that can arise from realistic differential migration.

In this work, we have achieved a more realistic distribution of pair-up radii via explicit integration of migration torques com-

puted for self-consistent SG AGN disk models modeled with the software pAGN (Gangardt et al. 2024). We identify migration traps as any location in the disk where the torque changes sign from positive to negative. Our key findings are as follows:

- The radial distribution of BBH pair-ups mainly follows the locations of migration traps, but with some spread and offsets due to differential migration. The overall pair-up radii profile is mostly influenced by the mass of the SMBH M_\bullet , specifically: (i) For $\log M_\bullet/M_\odot \lesssim 6.5$, BBH pair-ups occur broadly across the entire disk; (ii) for $\log M_\bullet/M_\odot \gtrsim 6.5$, pair-ups concentrate in the inner disk ($10^2 \lesssim R/R_S \lesssim 10^4$).
- The efficiency of BBH formation η is mildly influenced by the disk parameters, decreasing for increasing M_\bullet and α , and it is mostly influenced by the initial conditions of the migrating BHs. Indeed, BBHs are formed more efficiently for intermediate initial radii ($R_{\text{init}} \sim 10^2\text{--}10^5 R_S$) and large migrator mass ($m \gtrsim 20 M_\odot$).
- A large fraction of BBH pair-ups ($f_{\text{trap}} \gtrsim 80\%$ for first-generation BHs, $f_{\text{trap}} \sim 100\%$ for N -th generation BHs) occur in migration traps for low-to-intermediate SMBH masses $M_\bullet < M_\bullet^{\text{thr}}$. The threshold value is model dependent: $M_\bullet^{\text{thr}} \simeq 10^{7.4} M_\odot$ for the high-viscosity ($\alpha \geq 0.1$) in the B16 case, or $M_\bullet^{\text{thr}} \geq 10^{8.5} M_\odot$ for all other models. At larger SMBH masses, f_{trap} drops to lower values.
- For low SMBH masses ($\log M_\bullet/M_\odot < 6.0$), low viscosity ($\alpha = 0.01$) and B16 torque prescriptions, pair-up events pronouncedly pile up at $R \simeq 10^3 R_S$ as consequence of differential migration: sharp changes in the slope of $|\Gamma(R)|$ create a traffic-jam accumulation even in the absence of a trap.
- Estimating migration timescales as the inverse of the rate of angular momentum change ($t_{\text{migr}} \simeq L/\dot{L} = L/\Gamma$) tends to overestimate pair-up times in outer disk regions. True pair-up times (from simulations) are shorter due to differential migration and movement into regions of stronger torque.

Our results demonstrate that, although BBH formation predominantly happens in migration traps for a large portion of the explored parameter space, formation in the bulk is non-negligible.

By integrating the migration trajectories of individual BHs and applying physically motivated pair-up criteria, we recover a distribution of pair-up radii that reflects the full complexity of orbital evolution in AGN disks. This highlights the limitations of using predetermined locations to represent the entire BBH formation process.

Incorporating realistic pair-up radius distributions into semi-analytic population models will improve the physical accuracy of merger rate predictions. It allows for a more sensible representation of where and when binaries are likely to form, and it captures the contribution of dynamically overtaking encounters which is entirely missed when assuming a static trap location.

Data availability

Supplementary material, data, and code used in this study are available via Zenodo at [this link](#).

Acknowledgements. We thank the anonymous referee for their constructive and insightful comments that improved this manuscript. We thank Kees Dullemond, Dominika Wylezalek, Ralf Klessen, Evgeni Grishin, Nick Stone, Alessandro A. Trani, K.E. Saavik Ford, Barry McKernan, Shawn Ray, Thomas Rometsch, Zoltan Haiman, Shmuel Gilbaum, and Lumen Boco for useful discussions. MPV and MM acknowledge financial support from the European Research Council for the ERC Consolidator grant DEMOBLACK, under contract no. 770017. They also acknowledge financial support from the German Excellence Strategy via the Heidelberg Cluster of Excellence (EXC 2181 – 390900948)

STRUCTURES. The authors acknowledge support by the state of Baden-Württemberg through bwHPC and the German Research Foundation (DFG) through grants INST 35/1597-1 FUGG and INST 35/1503-1 FUGG. The public version of FASTCLUSTER is available via gitlab following [this link](#). We thank Daria Gangardt for developing the pAGN software, which is publicly available via github following [this link](#). This research made use of NUMPY (Harris et al. 2020), SCIPY (Virtanen et al. 2020), PANDAS (The pandas development Team 2024). Plots were produced using MATPLOTLIB (Hunter 2007).

References

- Abbott, R., Abbott, T. D., Acernese, F., et al. 2023a, *Phys. Rev. X*, **13**, 041039
- Abbott, R., Abbott, T. D., Acernese, F., et al. 2023b, *Phys. Rev. X*, **13**, 011048
- Bartos, I., Kocsis, B., Haiman, Z., & Márka, S. 2017, *ApJ*, **835**, 165
- Bellovary, J. M., Low, M. M., McKernan, B., & Saavik Ford, K. E. 2016, *ApJ*, **819**, L17
- Bryden, G., Chen, X., Lin, D. N. C., Nelson, R. P., & Papaloizou, J. C. B. 1999, *ApJ*, **514**, 344
- Cook, H. E., McKernan, B., Ford, K. E. S., et al. 2025, *ApJ*, **993**, 163
- Crida, A., Morbidelli, A., & Masset, F. 2006, *Icarus*, **181**, 587
- DeLaurentiis, S., Epstein-Martin, M., & Haiman, Z. 2023, *MNRAS*, **523**, 1126
- Delfavero, V., Ford, K. E. S., McKernan, B., et al. 2025, *ApJ*, **989**, 67
- Dittmann, A. J., Dempsey, A. M., & Li, H. 2025, *ApJ*, **990**, 137
- Duffell, P. C. 2015, *ApJ*, **807**, L11
- Epstein-Martin, M., Stone, N., & Becker, J. 2025, ArXiv e-prints [arXiv:2510.12895]
- Fung, J., Shi, J.-M., & Chiang, E. 2014, *ApJ*, **782**, 88
- Gangardt, D., Trani, A. A., Bonnerot, C., & Gerosa, D. 2024, *MNRAS*, **530**, 3689
- Gilbaum, S., Grishin, E., Stone, N. C., & Mandel, I. 2025, *ApJ*, **982**, L13
- Goodman, J. 2003, *MNRAS*, **339**, 937
- Greene, J. E., & Ho, L. C. 2007, *ApJ*, **667**, 131
- Grishin, E., Gilbaum, S., & Stone, N. C. 2024, *MNRAS*, **530**, 2114
- Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Nature*, **585**, 357
- Hawley, J. F., Guan, X., & Krolik, J. H. 2011, *ApJ*, **738**, 84
- Hill, G. W. 1878, *Am. J. Math.*, **1**, 5
- Hohmann, W. 1960, *The Attainability of Heavenly Bodies*, NASA Technical Translation (National Aeronautics and Space Administration)
- Hunter, J. D. 2007, *Comput. Sci. Eng.*, **9**, 90
- Iorio, G., Mapelli, M., Costa, G., et al. 2023, *MNRAS*, **524**, 426
- Ishibashi, W., & Gröbner, M. 2020, *A&A*, **639**, A108
- Ishibashi, W., & Gröbner, M. 2024, *MNRAS*, **529**, 883
- Jiménez, M. A., & Masset, F. S. 2017, *MNRAS*, **471**, 4917
- Kanagawa, K. D., Tanaka, H., & Szuszkiewicz, E. 2018, *ApJ*, **861**, 140
- King, A. R., Lubow, S. H., Ogilvie, G. I., & Pringle, J. E. 2005, *MNRAS*, **363**, 49
- Kley, W., & Nelson, R. P. 2012, *ARA&A*, **50**, 211
- Kocsis, B., Yunes, N., & Loeb, A. 2011, *Phys. Rev. D*, **84**, 024032
- Li, J., Dempsey, A. M., Li, H., Lai, D., & Li, S. 2023, *ApJ*, **944**, L42
- Lin, D. N. C., & Papaloizou, J. C. B. 1993, *Protostars and Planets III* (Tucson: University of Arizona Press), 749
- Lyra, W., Paardekooper, S.-J., & Low, M.-M. 2010, *ApJ*, **715**, L68
- Maggiore, M. 2018, *Gravitational Waves: Volume 2: Astrophysics and Cosmology*, Gravitational Waves (Oxford University Press)
- Malik, M., Meru, F., Mayer, L., & Meyer, M. 2015, *ApJ*, **802**, 56
- Mapelli, M., Dall’Amico, M., Bouffanais, Y., et al. 2021, *MNRAS*, **505**, 339
- Mapelli, M., Bouffanais, Y., Santoliquido, F., Arca Sedda, M., & Artale, M. C. 2022, *MNRAS*, **511**, 5797
- Martin, R. G., Nixon, C. J., Pringle, J. E., & Livio, M. 2019, *New Astron.*, **70**, 7
- Masset, F. S. 2017, *MNRAS*, **472**, 4204
- McKernan, B., Ford, K. E. S., Lyra, W., & Perets, H. B. 2012, *MNRAS*, **425**, 460
- McKernan, B., Ford, K. E. S., Kocsis, B., Lyra, W., & Winter, L. M. 2014, *MNRAS*, **441**, 900
- McKernan, B., Ford, K. E. S., O’Shaughnessy, R., & Wysocki, D. 2020, *MNRAS*, **494**, 1203
- McKernan, B., Ford, K. E. S., Cook, H. E., et al. 2025, *ApJ*, **990**, 217
- Moncrieff, J. W. N., Grishin, E., Trani, A. A., Panther, F. H., & Pietrosanti, O. 2026, *MNRAS*, **545**, staf2217
- Paardekooper, S. J., Baruteau, C., Crida, A., & Kley, W. 2010, *MNRAS*, **401**, 1950
- Penna, R. F., Sądowski, A., Kulkarni, A. K., & Narayan, R. 2012, *MNRAS*, **428**, 2255
- Peters, P. C. 1964, *Phys. Rev.*, **136**, 1224
- Petrovich, C., & Rafikov, R. R. 2012, *ApJ*, **758**, 33
- Rowan, C., Boekholt, T., Kocsis, B., & Haiman, Z. 2023, *MNRAS*, **524**, 2770
- Rowan, C., Whitehead, H., Boekholt, T., Kocsis, B., & Haiman, Z. 2024, *MNRAS*, **527**, 10448
- Rowan, C., Whitehead, H., Fabj, G., et al. 2025a, *MNRAS*, **543**, 132
- Rowan, C., Whitehead, H., Fabj, G., et al. 2025b, *MNRAS*, **539**, 1501
- Rozner, M., Generozov, A., & Perets, H. B. 2024, *ApJ*, **968**, 80
- Rozner, M., Generozov, A., & Perets, H. B. 2023, *MNRAS*, **521**, 866
- Samsing, J., Bartos, I., D’Orazio, D. J., et al. 2022, *Nature*, **603**, 237
- Secunda, A., Bellovary, J., Mac Low, M.-M., et al. 2020, *ApJ*, **903**, 133
- Shakura, N. I., & Sunyaev, R. A. 1973, *A&A*, **24**, 337
- Sirko, E., & Goodman, J. 2003, *MNRAS*, **341**, 501
- Stone, N. C., Metzger, B. D., & Haiman, Z. 2017, *MNRAS*, **464**, 946
- Tagawa, H., Haiman, Z., Bartos, I., & Kocsis, B. 2020a, *ApJ*, **899**, 26
- Tagawa, H., Haiman, Z., & Kocsis, B. 2020b, *ApJ*, **898**, 25
- Tagawa, H., Kimura, S. S., Haiman, Z., Perna, R., & Bartos, I. 2023, *ApJ*, **950**, 13
- The pandas development Team 2024, <https://doi.org/10.5281/zenodo.10697587>
- Tiede, C., Zrake, J., MacFadyen, A., & Haiman, Z. 2020, *ApJ*, **900**, 43
- Tornamenti, S., Mapelli, M., Périgois, C., et al. 2024, *A&A*, **688**, A148
- Trinca, A., Schneider, R., Valiante, R., et al. 2022, *MNRAS*, **511**, 616
- Vaccaro, M. P., Mapelli, M., Périgois, C., et al. 2024, *A&A*, **685**, A51
- Velasco Romero, D. A., & Masset, F. S. 2020, *MNRAS*, **495**, 2063
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *Nat. Meth.*, **17**, 261
- Wang, Y.-H., McKernan, B., Ford, S., et al. 2021, *ApJ*, **923**, L23
- Wang, Y., Zhu, Z., & Lin, D. N. C. 2024, *MNRAS*, **528**, 4958
- Whitehead, H., Rowan, C., Boekholt, T., & Kocsis, B. 2024a, *MNRAS*, **533**, 1766
- Whitehead, H., Rowan, C., Boekholt, T., & Kocsis, B. 2024b, *MNRAS*, **531**, 4656
- Whitehead, H., Rowan, C., & Kocsis, B. 2025, *MNRAS*, **543**, 3768
- Xue, L., Tagawa, H., Haiman, Z., & Bartos, I. 2025, *Phys. Rev. D*, **112**, 063034
- Yang, Y., Bartos, I., Gayathri, V., et al. 2019a, *Phys. Rev. Lett.*, **123**, 181101
- Yang, Y., Bartos, I., Haiman, Z., et al. 2019b, *ApJ*, **876**, 122
- Zevin, M., & Holz, D. E. 2022, *ApJ*, **935**, L20




Fig. A.1. Radial profiles of surface density Σ_g (top panel) and aspect ratio $h = H/R$ (bottom panel) of $\alpha = 0.1$ disks, under different numerical treatments. The thin cyan line shows the original SG profile as published in [Sirko & Goodman 2003](#), valid only for $\log M_*/M_\odot = 8.0$ and $\alpha = 0.01$. Thicker colorful lines are extensions of the model to different SMBH masses (orange: $\log M_*/M_\odot = 6.0$; purple: 7.0; black: 8.0). Different line-styles distinguish the numerical method: solid lines are self-consistent solutions of the SG equations computed with pAGN, while dashed and dotted lines represent power-law scaling approximations from [Vaccaro et al. \(2024\)](#) and [Grishin et al. \(2024\)](#) respectively. The dot-dashed gray line is a [Shakura & Sunyaev \(1973\)](#) disk for comparison. Solid lines should be regarded as the most physically consistent profiles within this framework.

Appendix A: Improved disk modeling and comparison to previous work

A proper and realistic description of the properties of AGN disks is crucial for modeling compact object dynamics in such environments. In this work, we model AGN disks using the analytical tool pAGN, developed by [Gangardt et al. \(2024\)](#), to numerically solve [Sirko & Goodman \(2003, SG\)](#) equations of disk structure. A similar approach has recently been adopted in most semi-analytical models for the AGN channel (e.g. [Vaccaro et al. 2024](#); [Cook et al. 2025](#); [Delfavero et al. 2025](#)). This provides smooth profiles for key quantities for any combination of system parameters in a manner fully consistent with the physical description of SG. The resulting profiles for selected SMBH masses M_* are shown in Figure A.1 and Figure A.2 (hereafter A.1 and A.1) for high ($\alpha = 0.1$) and low ($\alpha = 0.01$) viscosity respectively.

In this appendix, we compare our current prescription with models of α -disks previously used in the literature. An α -disk




Fig. A.2. Same as Figure A.1, but for $\alpha = 0.01$.

is a model in which gas turbulence is assumed to be the cause of disk viscosity, characterized by the viscosity coefficient $\alpha \in [0, 1]$. The prototypical α -disk is the one by [Shakura & Sunyaev \(1973\)](#), where crucial quantities such as the gas surface density Σ_g and aspect ratio h are modeled as power laws. We show the resulting profile in A.1 and A.1, using the [Kocsis et al. \(2011, eqs. 32–33\)](#) treatment and assuming $M_*/M_{*,\text{Edd}} = 0.1$. The obvious problem with this disk model is that becomes unphysical in the outer regions, and it is hence only valid for $R \lesssim 10^3 R_S$. Moreover, such disks are viscously, thermally, and convectively unstable to perturbations, meaning that they can easily fragment into smaller sub-clouds ([Kocsis et al. 2011](#)).

[Sirko & Goodman \(2003\)](#) extend the [Shakura & Sunyaev \(1973\)](#) model by assuming that star formation in the outskirts of the disk acts as stabilizing mechanism against fragmentation. Moreover, the SG model is also broadly consistent with observed quasar spectral energy distributions, as it more accurately reproduces the flat UV–optical slopes and IR excesses characteristic of luminous quasars. The full model consists of 16 coupled equations describing the radial structure of the disk, which were originally solved self-consistently for a fiducial set of parameters: $\log M_*/M_\odot = 8.0$, $\alpha = 0.01$, and $M_*/M_{*,\text{Edd}} = 0.1$. Before the publication of the pAGN tool, various approximations were required to generalize this solution to other physical regimes. In the following, we illustrate two examples.

In [Vaccaro et al. \(2024, eqs. 1–2\)](#), we fit the original profile with broken power laws, and combine it with a first-order rescaling in M_* and α . As shown in A.1 and A.1, this provides quite a crude modeling of the relevant quantities, as it neglects the




Fig. A.3. Two-dimensional histograms showing the distribution of BBH pair-up locations in the **K18** model in the radius–mass plane (R vs M_*), for different combinations of torque prescription (top: **G24**; bottom: **B16**) and disk viscosity α (left: 0.01; right: 0.1), for a **SG** disk model. The color map indicates the number of pair-ups in each bin, N_{pair} . Pink dots mark the locations of migration traps in the **K18** model, while red crosses mark the Type I trap positions (as in Figure 2). Pink contours show the values of the normalized contrast parameter K_{norm} .

complex interdependency between parameters in the full system of equations. The approximation breaks down most noticeably at high α and high M_* , where it severely overestimates the surface density Σ_g in the inner regions of the disk.

Another approach is the one by [Grishin et al. \(2024, eqs. 3–11\)](#) who reengineer an analytical model which is similar in spirit to **SG**, but neglecting some of the physics and using simplified analytical expressions. The resulting profiles, also shown in [A.1](#) and [A.1](#) for $M_*/M_{\text{Edd}} = 0.1$, are qualitatively similar to the **SG** ones computed with **pAGN** (meaning that they feature an inner region where the density increases and the thickness decreases for growing radii, and an outer region where the opposite happens), but the transition radii are systematically one order of magnitude smaller than in the more accurate **SG** paradigm. Moreover, they also overestimate the density and underestimate the aspect ratio in the inner regions ($R \lesssim 10^3 R_S$).

This highlights the importance of solving the full set of equations with **pAGN**, as the dependence of the disk structure on α and

M_* is highly nonlinear and cannot be captured by simple rescaling or analytical fits, when constructing models of BH dynamics in AGN disks.

Appendix B: Type II migration

The classical picture of Type II migration assumes that, once a gap is opened, the embedded object becomes locked to the viscous inflow of the disk with migration timescale $t_{\text{type II}} = t_{\text{visc}}$, as assumed in subsection 2.2. This disk-dominated view is conservative: it assumes the gap to be completely cleared and gas inflow across the gap to be suppressed. Recent hydrodynamical simulations have shown that this picture is oversimplified: in reality, gap opening is often partial and gas can leak through the gap edges, leading to migration rates that are faster than the purely viscous timescale (e.g. [Malik et al. 2015](#)).

[Kanagawa et al. \(2018\)](#) (hereafter **K18**) proposed a modified prescription that interpolates between Type I and Type II

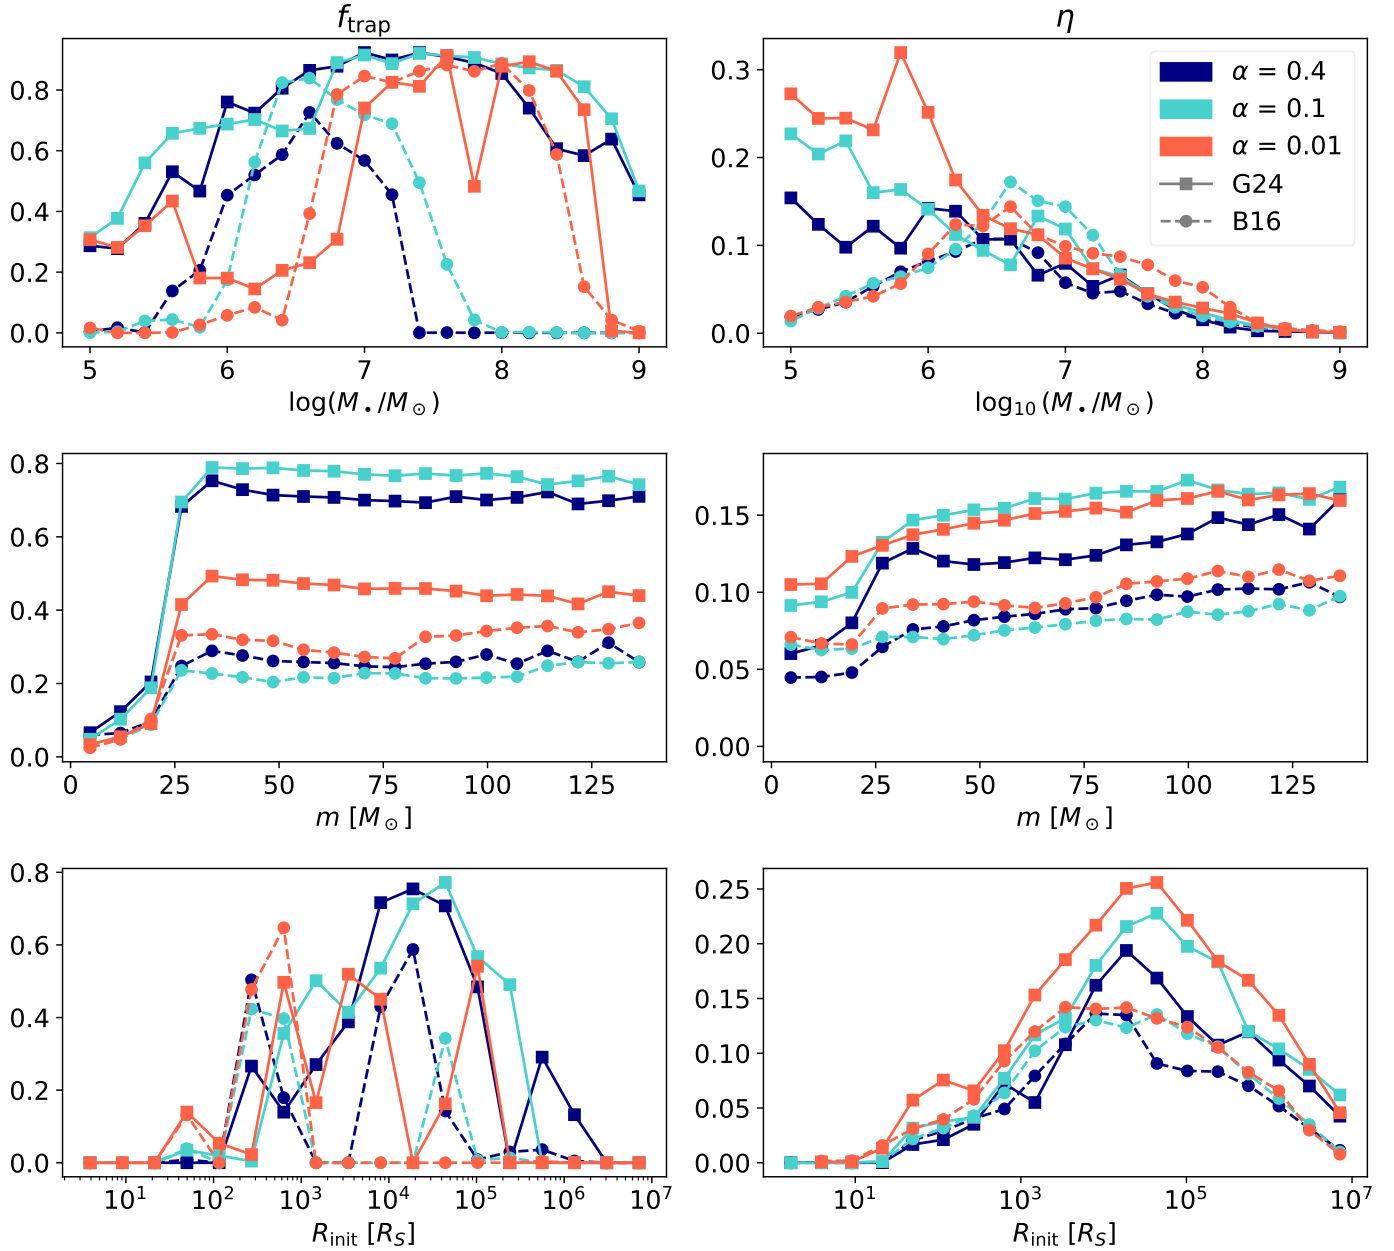


Fig. B.1. Fraction of pair-ups occurring within migration traps (left panels) and pair-up efficiency (right panels) as functions of SMBH mass M , (top row), migrator mass m (middle row), and initial radius R_{init} (bottom row) in the K18 prescription. Color encodes the viscosity parameter α (orange: 0.01; teal: 0.1; navy blue: 0.4), and the two torque prescriptions are distinguished by marker shape and line style (squares and solid lines: G24; circles and dashed lines: B16).

regimes, so that migration is well described as a modified Type I regime with reduced local density

$$\Sigma'_g(R) = \frac{\Sigma_g(R)}{1 + 0.04 K}, \quad (\text{B.1})$$

where the contrast parameter K is defined as in eq. 13. For low-mass migrators ($K \lesssim 1$, $\Sigma'_g \gtrsim 0.96 \Sigma_g$), the gap is shallow and standard Type I migration (eq. 1) remains valid. As the migrator mass increases, the contrast parameter K grows quadratically. The full transition to a Type II regime happens for $K \gtrsim 20$ (K18), as the disk surface density at the orbit is significantly reduced to $\Sigma'_g \lesssim 0.55 \Sigma_g$.

We calculate corotational gas torques in eqs. 3, 4, 5, 6, 7, and 12 using the reduced local density Σ'_g . In this regime,

both the normalization of the torque Γ_0 (eq. 12), and the shape of the torque profiles (via the parameter a in eq. 4) are significantly modified. We include an exponential decay term which scales with the contrast factor K , as in K18. We leave the Lindblad torque (Γ_L , eq. 8) unperturbed.

The effect of gap opening of Lindblad torques is quite complex, as it depends on the exact profile of the gap (e.g. Kley & Nelson 2012; Petrovich & Rafikov 2012). A more comprehensive simulation would include modeling the gap's density profile and computing Lindblad torques self-consistently. However, in the high-viscosity regime in which we are operating, we do not expect any nonlinear behavior that would strongly affect our results (Duffell 2015).

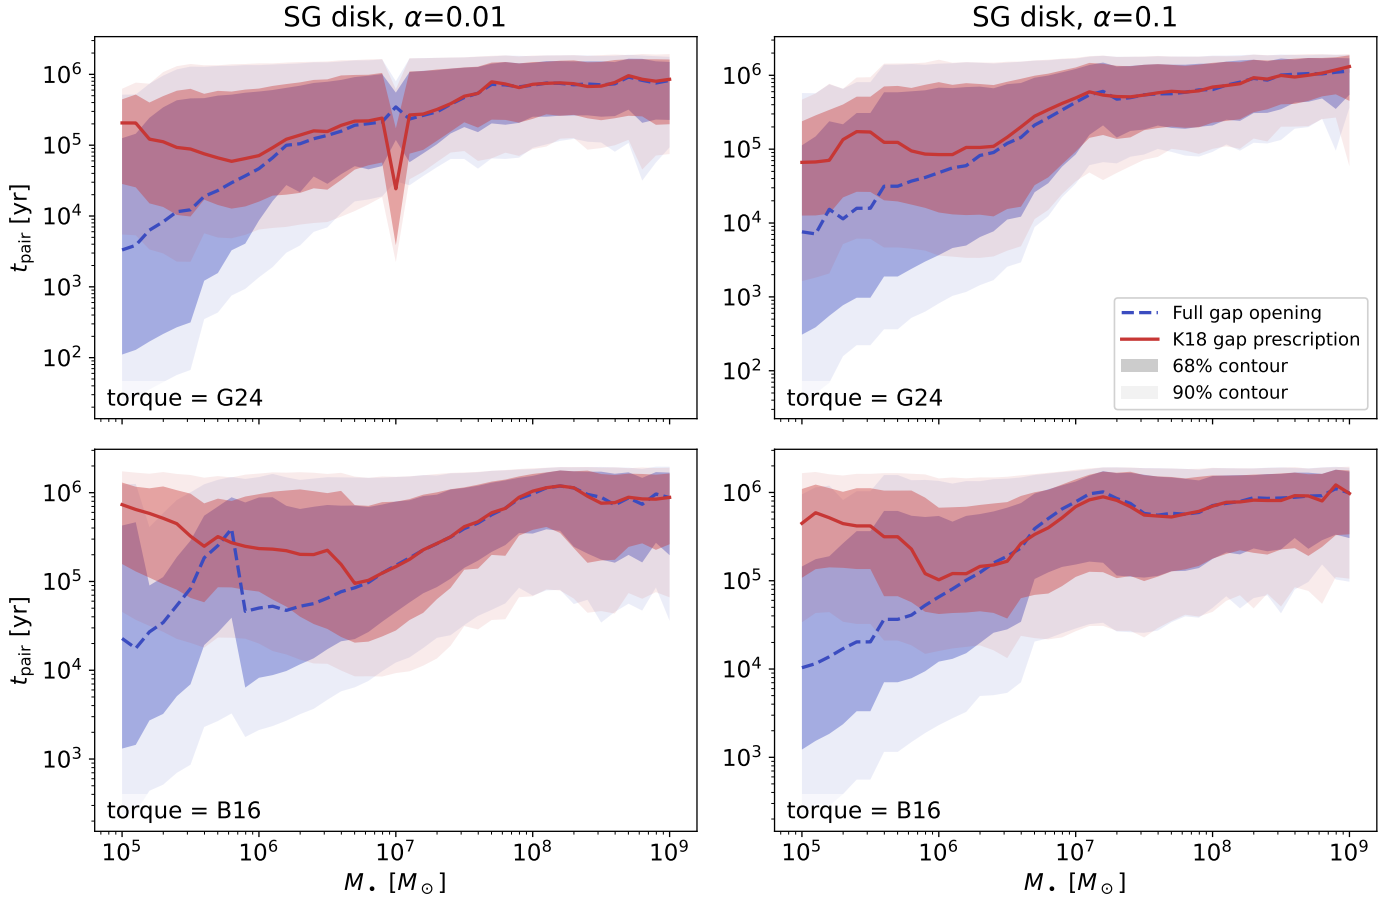


Fig. B.2. Pair-up times t_{pair} as a function of SMBH mass in SG AGN disks, for different combinations of the viscosity parameter α (left: 0.01; right: 0.1) and torque prescriptions (top: G24; bottom: B16). Each panel shows a median line and contours enclosing 68% and 90% of the pair-up time distribution, based on our simulations. Results are shown for two different Type II migration prescriptions, indicated by color and line style. In blue, dashed lines: full gap opening, as in the main text (eq. 15). In red, continuous lines: partial gap opening, as in this appendix (K18).

Explicitly, our torque equations read

$$\Gamma' = \left[\left(\frac{0.46 - 0.96a' + 1.8b}{\gamma} \right) \Gamma'_0 + \Gamma'_{\text{th}} \right] e^{-K/20} + \Gamma_L \quad (\text{B.2})$$

in the G24 case and

$$\Gamma' = \frac{\Gamma'_{\text{ad}}(\Theta')^2 + \Gamma'_{\text{iso}}}{(\Theta' + 1)^2} e^{-K/20} \quad (\text{B.3})$$

in the B16 case, where $\Gamma'_i = \Gamma_i(\Sigma'_g)$ and $a' = -\partial \log \Sigma'_g / \partial \log R$.

An important implication is that migration traps are no longer determined solely by the background disk structure, as the torques $\Gamma' = \Gamma(\Sigma'_g)$ are now K -dependent and, consequently, m -dependent. This mass dependence adds an additional layer of complexity to the dynamics of BHs in AGN disks. In the following, the trap locations under the K18 prescription are defined as the radii where the total migration torque Γ' acting on a migrator of average mass ($m \approx 13.5 M_\odot$) changes sign from positive to negative.

While the gap contrast naturally depends on the properties of the migrator, a disk-specific contrast parameter K_{norm} can be defined by normalizing over the migrator mass:

$$K_{\text{norm}} = \frac{K}{m^2} = \alpha^{-1} h^{-5} \left(\frac{M_*}{M_\odot} \right)^{-2} \quad (\text{B.4})$$

We recompute the binary formation catalog using the K18 prescription for gap opening. The resulting pair-up distributions

are shown in Figure A.3. Compared to the results presented in subsection 3, the locations of migration traps differ substantially for low-mass SMBHs ($M_* \lesssim 10^7 M_\odot$), because the modified torques tend to shift the zero-torque radii away from regions of high contrast K_{norm} .

In the G24 scenario, the peaks in the pair-up radius distributions are generally located in proximity to the new migration traps, with the exception of systems with $M_* \lesssim 10^6 M_\odot$ which present traffic-jam accumulations of pair-up events in regions of low contrast ($K_{\text{norm}} \lesssim 1$), even in the absence of traps.

In the B16 framework, the Lindblad torque is neglected (see eqs. 2 and B.3). Hence, the migration torque is highly suppressed in areas of high contrast ($K_{\text{norm}} \gtrsim 1$) and almost no pair-up events happen at radii $R \in [10^2, 10^5] R_g$ for $M_* \leq 10^{5.8} M_\odot$. Moreover, disks with SMBH masses up to $10^{6.5} M_\odot$ present prominent traffic-jams in the outer area of the disks.

Figure B.1 shows the resulting pair-up efficiency, η , and trap fraction, f_{trap} , obtained with the K18 model.

The pair-up efficiency, η , is generally lower than in the fiducial runs, peaking at roughly 0.3 for $M_* \lesssim 10^6 M_\odot$ in the low-viscosity G24 case, while retaining similar trends in SMBH mass M_* and in initial radius R_{init} to those observed in subsection 3. On the other hand, the dependency on migrator mass m is flatter than in the fiducial case especially for disks with low viscosity ($\alpha = 0.01$), implying that partial gap opening (K18) moderately suppresses the pair-up efficiency of larger mass migrators (with $m \gtrsim 20 M_\odot$).

The fraction of pair-ups occurring within migration traps, f_{trap} , retains a similar behavior than in Figure 3 for the high-viscosity G24 case, remaining large (> 0.6) for most of the parameter space. In the low-viscosity G24 case, f_{trap} is large for $M_{\bullet} \in [10^7, 10^{8.7}] M_{\odot}$ and $\lesssim 0.3$ otherwise. Contrarily, in the B16 scenario, traffic-jam pair-ups are more relevant, especially at low SMBH mass. Therefore, $f_{\text{trap}} < 0.2$ with the exception of a certain α -dependent window: $\log M_{\bullet}/M_{\odot} \in [6.5, 8.6]$ for $\alpha = 0.01$, $\log M_{\bullet}/M_{\odot} \in [6.0, 7.6]$ for $\alpha = 0.1$, or $\log M_{\bullet}/M_{\odot} \in [5.8, 7.3]$ for $\alpha = 0.4$. It retains a similar trend in both migrator mass m and initial radius R_{init} than discussed in subsection 3.4.

We compute pair-up times, t_{pair} , and evaluate how the K18 framework affects them. These pair-up times, shown via the 68% and 90% contour levels in Figure B.2, represent the elapsed time between the moment at which the single primary BH becomes embedded in the disk and the subsequent formation of a BBH. At low SMBH mass, $M_{\bullet} \lesssim 3 \times 10^6 M_{\odot}$, the normalized contrast parameter can reach larger values, $K_{\text{norm}} \gtrsim 1$, thus significantly reducing the local gas density Σ' and slowing down the pair-up process, therefore increasing t_{pair} by almost two orders of magnitude for the lowest SMBH mass explored, $M_{\bullet} = 10^5 M_{\odot}$. In contrast, for $M_{\bullet} \gtrsim 3 \times 10^6 M_{\odot}$, pair-up times are unaffected by the adopted Type II prescription.

Overall, the adoption of the K18 prescription has different consequences depending on the torque model adopted.

In the G24 case, the position of traps is shifted toward regions of low contrast K_{norm} . For $M_{\bullet} \lesssim 10^6 (10^7) M_{\odot}$ in the high-viscosity (low-viscosity) case, pair-up events happen in traffic jams in valleys of low contrast K_{norm} . For larger SMBH masses, both the pair-up efficiency and the relevance of migration traps are mostly unaffected.

In the B16 case, the pair-up process in disks with low SMBH mass ($\log M_{\bullet}/M_{\odot} \lesssim 6$) is strongly suppressed, and pair-up events often happen in the outskirts of the disk.

Consequently, the self-consistent treatment of gap opening lowers the pair-up efficiency of high-mass migrators and weakens the correlation between migration traps and binary formation sites.